UNIVERSITY OF GOTHENBURG
SCHOOL OF BUSINESS, ECONOMICS AND LAW

An empirical evaluation of Value at Risk

Master Thesis - Industrial and financial management
University of Gothenburg,
School of Business, Economics & Law
January 2009

Instructor: Zia Mansouri
Authors: Martin Gustafsson 1982
Caroline Lundberg 1984

ABSTRACT

In light of the recent financial crisis, risk management has become a very current issue. One of the most intuitive and comprehendible risk measures is Value at Risk (VaR). VaR puts a monetary value on the risk that arises from holding an asset and is defined as the *“the worst loss over a target horizon with a given level of confidence”.* There are currently numerous techniques for calculating VaR on the market and no standard is set on how it should be done. This can make the field of VaR hard to overlook for someone not initiated in the world of econometrics.

In this paper we examine three basic widely used approaches used to calculate VaR. The approaches examined are the non Historical Simulation approach, the GARCH approach and the Moving Average approach. This paper has two main purposes, the first is to test the different approaches and compare them to each other in terms of accuracy. The second is to analyze the results and see if any conclusions can be drawn from the accuracy of the approach with respect to the return characteristics of the underlying assets. The accuracy of a VaR approach is tested by the number of VaR breaks that it produces i.e. the number of times that the observed asset return exceeds the predicted VaR. The number of VaR breaks is evaluated with the Kupiec test which defines an interval of VaR breaks in which the approach must perform to be accepted. By the term “*asset return characteristics”* we mean the statistical properties of the returns of the assets such as volatility, kurtosis and skewness.

The study is conducted on three fundamentally different assets, Brent crude oil, OMXs30 and Swedish three months treasury bills. Daily return data has been collected starting from January 1st 1987 to September 30th 2008 for all the assets. Observations spanning from 1987 to 1995 will be used as historical input for the approaches and observations spanning from 1996 to September 30th 2008 will be the comparative period in which the approaches performance will be measured.

The results of the study show that none of the three approaches are superior to the others and that more complex approaches do not guarantee more accurate results. Instead it seems that the characteristics of the asset returns in combination with the desired confidence level determine how well a certain approach performs on a certain asset. We show that the choice of VaR approach should be evaluated individually depending on the assets to which it is to be applied.

Keywords: Value at Risk, return characteristics, historical simulation, moving average, GARCH, normal distribution, Brent oil, OMXs30, Swedish treasury bills.

GLOSSARY OF MAIN TERMS

Autocorrelation: Is the correlation between the returns at different points in time.

Backtesting: The process of testing a trading strategy on prior time periods.

Fat tails: Tails of probability distributions that are larger than those of normal distribution.

GARCH approach to VaR: An approach for forecasting volatilities that assumes that volatility next period depends on lagged volatilities and lagged squared returns.

Historical simulation approach to VaR: An approach that estimates VaR from a profit and loss distribution simulated using historical returns data.

Kupiec test: Is a valuation method to evaluate VaR results.

Kurtosis: Measures the peakedness of a data sample. A high value of kurtosis means that more of the data's variance comes from extreme deviations.

Moving average approach to VaR: A parametric approach that assumes normal distributed returns and disregards the fact of volatility clustering.

Normal distribution: The Gaussian or bell-curve probability distribution.

Outlier: An extreme rare return observation.

Risk: The prospect of gain and loss. Risk is usually regarded as quantifiable.

Skewness: Tells us whether the data is symmetric or not.

Value at Risk (VaR): The maximum likely loss over some particular holding period at a particular level of confidence.

Volatility: The variability of a price, usually interpreted as its standard deviation.

Volatility clustering: High- volatility observations seems to be clustered with other high-volatility observations and the same is true for low-volatility observations that cluster with other low-volatility observations.

LIST OF FIGURES

TABLE OF CONTENTS

1. INTRODUCTION

This chapter gives a background to the subject treated in this study- Value at Risk. It further states the problems associated with the calculation of Value at Risk and the purpose of this paper. The chapter ends with the disposition of the paper.

1.1 WHAT IS VALUE AT RISK

To answer this question we can start off by asking ourselves the more fundamental question of what risk is. In economics there are many different types of risk such as legal risk, market risk, company specific risk and so on. But how do we define what risk is and how do we numerically measure its size? In finance, risk is often thought of as the volatility or standard deviation of an assets returns. But all volatility really tells us is how much the returns of an asset varies around its mean. This labels both positive and negative price movements as risk, although most investors would consider risk something negative and upward price movements as something positive (Jorion 2001).

It is not a very intuitive measure and can be hard to grasp for those not initiated in the world of finance. What the value at risk measurement does is that it puts a number, either a money value or a percent value of *"the worst loss over a target horizon with a given level of confidence"* (Jorion 2001:22). Value at risk or VaR as we will address it from here on thus consists of three components, a confidence level, a time horizon and a value. The time component tells us how far into the future we are looking, the further we look the larger the potential loss. The confidence level determines with which certainty the measurement is made, higher confidence levels means higher potential losses. Finally, the value component simply is the monetary value that we risk losing during the proposed time horizon given the proposed confidence level.

For example a $1000, one day, 95 percent confidence level VaR value for a stock means that during the next day we can be 95 percent certain that the value of our holdings in this particular stock will not decrease by more than $1 000.

An event where the return of an asset exceeds the estimated VaR measure is called a VaR break. The chosen confidence level of the VaR approach determines how many VaR breaks an approach should produce if performing well. The aim of a VaR approach is therefore not to produce as few VaR breaks as possible. An accurate VaR approach produces a number of VaR breaks as close as possible to the number of VaR breaks specified by the confidence level. Therefore if a VaR approach with 95% confidence is calculated it should produce VaR measures that are exceeded by the return of the asset five percent of the times it is applied. Thus if VaR predictions are made for 1000 individual days the estimated VaR should be exceeded 50 times. Any deviation from the expected number of VaR breaks regardless if it is negative or positive is a sign of inaccuracy or miss specification.

1.2 BACKGROUND OF VaR

Risk management is an important part for all financial institutions and also for companies that are exposed to risk. During the last decade there has been a revolution in risk management and VaR is one of the measures that have gotten a lot of attention. Holton (2003) writes that the roots of VaR however date back to as early as 1922, at which time The New York Stock Exchange imposed capital requirements on member firms.

Research on VaR did however not take off until 1952 when two researchers, Markowitz and Roy, almost simultaneously but independent of each other published quite similar methods of measuring VaR. According to Holton (2003), they were working on developing a means of selecting portfolios that would be able to optimize the reward for a given level of risk. He further points out that after that it took another 40 years until the VaR measurement began to be widely used among financial institutions and companies.

According to Fernandez (2003), past financial disasters such as the one in 1987 and the crises that followed led to a decision by the Basel Committee that all banks should keep enough cash to be able to cover potential losses in their trading portfolios over a ten-day horizon, 99 percent of the time. The amount of cash to be kept was to be calculated using VaR. Past crises has shown that enormous amounts of money can be lost over a day as a result of poor supervision and financial risk management. The VaR measure therefore got its breakthrough because of historical mistakes in risk management.

Today, the utilization of VaR is widely spread in financial institutions, however it is not as widespread in non financial firms. This can be explained by the fact that non financial firms do not usually forecast profits and losses on a daily basis. Nevertheless Mauro (1999) points out that VaR can be, and is used, in non financial firms that are affected by volatility in prices, particularly in a short time horizon.

An advantage of VaR is that it is a measure that can be applied to almost any asset. A disadvantage however, is that there is almost an infinite number of ways to calculate VaR, each of the approaches with its own advantages and disadvantages. This makes different VaR measurements hard to compare with each other if they have not been calculated using the same approach.

1.3 PROBLEM DISCUSSION

Value at Risk is a measure that is widespread within financial institutions where the importance of strict risk management has become vital. According to Jorion (2002) VaR has become the standard benchmark for measuring financial risk. Banks with large trading portfolios and institutions that deal with numerous sources of financial risk have been heading the use of risk management. Jorion (2001) writes that to put a number on the possible maximum loss under a specific time horizon has for many banks become an obligation. An example of this is Lesley Daniels Webster, former head of market risk at Chase Manhattan Bank, who saw it as a necessity. Every morning he received a 30-page report that summarized the VaR of the bank. The neat little report quantified the risk of all the trading positions of the bank. VaR is an important measure within risk management because it is an intuitive risk measure as it puts a monetary number on the risk exposure a company faces. According to Linsmeier & Pearson (1996) the strengths of VaR is that it provides a simpler and accurate overall measure of the market risk the company is taking. This is done without going into specific and complex details about the company's positions. This is a great advantage since risk assessments are most often produced for senior management with no or little experience with econometrics.

There does not, however, appear to be any consensus on how to calculate VaR. Today there seems to be as many ways to calculate VaR as there are practitioners using it. There are some general approaches to choose from, such as parametric, semi parametric or non parametric approaches. A parametric approach means that you use a parameter or a model to describe reality, often generalizing and making assumptions. A non parametric approach however may for example only look at historical data to make predictions about the future. Under each of these categories there are many different approaches that can be used to calculate VaR and each and every one of them can in their turn be applied in different ways under various assumptions or generalizations. In this study the historical simulation approach will be representing the group of non parametric approaches. Moving average and GARCH will be representing the parametric approaches. The GARCH approach considers volatility clustering i.e. the fact that days with high volatility are often clustered together. The approaches have been chosen because they are the most widely used approaches of calculating VaR.

There are several earlier studies on the subject, however most of them aim to develop new and more advanced ways of calculating VaR. Linsmeier and Pearson (1996) concludes that there is no simple answer to which VaR approach is the best. They all differ in their ability of capturing risk and further they differ in their ease of implementation and the ease of explanation to senior management. They state that the historical simulation approach is easy to implement but it does require a lot of historical data. It is also an approach that is easy to explain to senior management which is an essential part of the purpose of the VaR measure. In another study made by Goorbergh & Vlaar (1999) the most important characteristic of stock returns when using VaR is volatility clustering. This is a phenomenom that can be effectively modeled by means of GARCH.

The studies made in the past are not unanimous, they produce contradictory results. For example Cabedo & Moya (2003) finds that VaRs from an autoregressive moving average approach (ARMA) outperforms the GARCH. Costello et al. (2008) finds that the opposite is true, GARCH outperforms the ARMA and they suggest that the conclusion drawn by Cabedo & Moya is based on the assumption of normal distribution. The different approaches are applied on different assets and the approaches are often adapted to the assets to which they are applied. This indicate that the results from different studies can be hard to interpret and compare. We therefore think that it would be interesting to conduct a study by using the three most widely used approaches. They will all be applied on the same three assets to make the results more comparable.

The VaR will be estimated on a daily basis using three different confidence levels, 95%, 99% and 99.9% to see how the approaches perform at different confidence levels on different assets. Jorion

(2001) and many others with him uses these confidence levels when calculating VaR. The different levels fit different purposes and depends on the management's relation to risk. Choosing a higher level of confidence will result in a higher VaR. The difference between a confidence level of 99 or 99.9 percent might not seem big but has a large effect on the VaR assessment.

Different characteristics of the assets returns, i.e. the statistical properties of the assets returns such as volatility, kurtosis and skewness, might affect the calculated VaR making some approaches more preferable with specific assets. This is especially true with the parametric approaches that assume the returns to be normally distributed. This leads to the problem that this study is focused on, the connection between an assets return characteristics and its effect on the calculated VaR. Three different assets have been chosen; Brent crude oil, Swedish stock market index (OMXs30) and Swedish three month treasury bills (STB3M). The motives for choosing these assets are that they are fundamentally different assets with fundamentally different characteristics. The reason why this is so important in this study is that it provides a better prerequisite to examine how the return characteristics affect VaR approaches.

VaR can be applied on any asset that has a measurable return. The assets used in this study all have in common that their returns can easily be measured, their characteristics do however differ. Brent oil has many uses one being petrol production and another is heating purposes. Oil companies sitting on large reserves are very sensitive to changes in oil prices making VaR a suitable risk measure for assessing potential future losses. The same goes for anyone trading with oil regardless of them being buyers or sellers. The oil market differs from the stock market in the sense of buyers. Oil is traded in huge amounts by large oil companies in contrast to stocks that can be traded in small quantities by small investors. Stocks are perhaps the most common asset used in VaR calculations. The risks connected to stock returns are quite easy to understand. Financial institutions invest in huge portfolios consisting of various stocks with varying risks. The market risk they face must be evaluated and according to Jorion (2002) this is best done by using VaR. When it comes to treasury bills it should represent a stable and secure investment instrument. Fluctuating interest rates however affects the market value of the bills. This means that the market value of the bills can vary during their holding time which makes VaR a useful measure in order to estimate potential losses.

- How do the different characteristics of the returns of the underlying assets affect VaR?

Another important aspect is the complexity of the VaR approaches themselves. We will take a closer look on the performance difference between more complex and simpler approaches to calculating VaR. This will result in an evaluation of whether the extra effort put in to the use of the more complex approaches is worth the while in terms of accuracy. The accuracy of the approaches will be measured in terms of the number of VaR breaks they produce

- Which of the approaches gives the most accurate VaR measure?

1.4 PURPOSE

The purpose of this study is to compare three different approaches to calculating VaR namely the Historical Simulation approach, the Moving Average approach and the GARCH approach. The approaches will be applied on the three different assets with the confidence levels 95%, 99% and 99.9%. The performance and accuracy of the approaches will then be analyzed with respect to asset return characteristics and complexity of the approach.

1.5 DISPOSITION

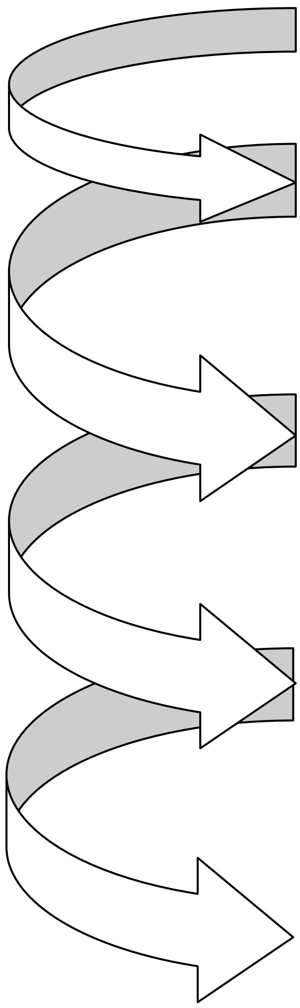

Figure 1.5 Disposition of this paper.

Introduction. The first chapter presents the subject examined in this study and gives an introduction to VaR as a measure in risk management and how it can be used.

Theory. In the upcoming chapter, the development and functions of Value at Risk are described along with the chosen methods of calculating Value at Risk. It is placed before the method chapter in order to give a better understanding for the approaches used making the calculations more understandable.

Method. The third chapter deals with the method used in the study. First, the analytical approach is described and then we move on to the methods used to calculate the Value at Risk. The data of the assets, which Value at Risk are calculated on, is also presented and their return characteristics are illustrated and described.

Results & Analysis. In the fourth chapter, the results are presented in form of the produced Value at Risk estimations. Back testing is made in order to see how the methods have performed. It also contains the analysis where the results are analyzed to see how the characteristics of the asset returns affect the calculated Value at Risk

Conclusions. In the fifth chapter, we present the conclusions that we have drawn from the calculations and discuss further research.

2. THEORY

In this chapter we will present the risk measure Value at risk and the different VaR approaches that we will be using in this thesis along with the mathematical models that they are based on. The underlying assets will also be presented.

2.1 VALUE AT RISK - VaR

As mentioned earlier the definition of VaR by Jorion (2001:22) is: *"VaR summarizes the worst loss over a target horizon with a given level of confidence"*. The most common approach is the parametric under which the mathematics of VaR can be described by the function below.

$$VaR = \sigma \times \mathrm{C} \times \sqrt{t} \times \$$$

The daily standard deviation times the number of standard deviations that corresponds to the selected confidence level, C, times the square root of the time horizon times the monetary size of the investment results in the VaR. In this thesis we will be calculating one day VaR estimates so the time component can be excluded. Also, the monetary size of the investment is just there to put a monetary figure on the risk and can also be excluded. What we have left is the standard deviation of historical returns and the confidence level.

If choosing a lower confidence level, as 95 percent or 90 percent, the VaR will decrease. Different levels will fit different firms and purposes and will be chosen according to the management's relation to risk. The more risk averse the firm is the higher confidence level will be selected. However Dowd (1998) claims that VaR users and than in first hand banks will in most cases prefer the lower confidence value in order to decrease the capital needed to cover the potential future losses.

In order to provide an overview of the measure VaR, a simple illustration is given. Suppose that the oil price today is 100 USD per barrel and that the daily standard deviation (σ) is 20 USD. Companies that buy large amounts of oil might want to know how much, given a certain confidence level, they can possibly lose when buying the oil today compared to tomorrow. If the chosen confidence interval is 99 percent, that would mean that one day out of hundred the loss will be greater than the calculated VaR. This is true when the price change is normal distributed around the average price change.

Value due to an increase in the oil price: $100 + 2.33 \times \sigma = 146.6$ USD
Value due to an decrease in the oil price: $100 + 2.33 \times \sigma = 53.4$ USD

This means that with 99 percent chance the loss will not be greater than 100-53.4 = 46.6 USD which is the VaR for a confidence level of 99 percent.

Often some assumption are made in order to calculate the VaR and if we suppose that the daily changes in oil prices has a density function this function is most often assumed to be represented by the normal distribution. The assumption has the benefit of making the VaR estimations much simpler. However it has some disadvantages. The price changes do not always fit the normal distribution curve and when more observations are found in the tails, normal-based VaR will understate the losses that can occur. A solution could be to use another distribution that regards fatter tails. The fatter tails that for example comes with the t-distribution makes high losses more common according to Dowd (1998) and therefore it also gives higher VaR.

When using VaR, the assets have to be valued at their market value. This will not be a problem according to Penza & Bansal (2001) with assets traded in a liquid market, where prices easily can be derived. This is called marking-to-market and Dowd (1998) defines the term as the practice of valuing and frequently revaluing positions in marketable securities by means of their current market prices. Financial assets are often traded on a daily basis and therefore, it is easier to obtain their current value. For some nonfinancial assets, the price may have to be estimated, which make the calculations more uncertain. The assets used in this study causes no problem in this area, however one can claim that the pricing does differ among them. Some claim to be able to predict the future price of a specific stock because of available information concerning the stock, while future prices on oil may be harder to predict. Reasons for this can be the strong influence OPEC has on the price of oil by changing the relation of demand and supply. Sadeghi & Shavvalpour (2006) call attention to the need of VaR as a tool for quantifying market risk within oil markets. Future asset prices and reasons for them are a complex area however as stated before all the assets used in this study are traded daily, making it no problem to derive the needed market prices.

2.2 THE NEED FOR VaR

Holton (2002) puts the breakthrough for VaR in the 1970s and 1980s. During this time many changes took place and diverse financial crises became a fact. The Bretton Woods system collapsed in 1971 making the exchange rates flow. Because of the oil crisis caused by OPEC, oil prices went through the roof by going from USD 2 to USD 35. Along with the crisis also came financial innovations such as the proliferation of leverage. Before this time, the avenues for compounding risk were limited according to Holton (2002). With new instruments and new forms of transactions, new leverage ways opened up. Holton also brings up that when this happened, trading organizations sought new ways to manage risk taking. During this time, banks started to create a measure as a result of more complex risk management. The risks had to be aggregated so when facing this problem, the wanted solution was a measure that could provide the companies with a better understanding for the risk exposure. It was JP Morgan who became the most successful. They created the RiskMetrics system which is mentioned by Jorion (2001). It revealed risk measures for 300 financial instruments and the data characterize a variance – covariance matrix of risk and correlation measures that progress through time. The development of the measure took a long time and was finished in 1990. Four years later it was made free to the public something that created big attention. The use of VaR system increased enormously and a positive effect on the companies risk management were found. The important contribution of RiskMetrics was that it publicized VaR to a wide audience.

An important landmark in risk management that made VaR known to the mass was the Basel Accord. It was concluded in 1988 and fully implemented in the G-10 countries[1] in 1992. This is mentioned by Saunders (1999) and it was a regulation on commercial banks to provide a more secure system through minimum capital requirements for banks' markets risk.

2.3 VaR APPROACHES

The three different approaches used in this study will be presented below. They all have their own advantages and disadvantages which affects the VaR values that they produce. The assumptions made by the approaches deals with the return characteristics in different ways. The effects of this will be further discussed in the analysis.

2.3.1 HISTORICAL SIMULATION APPROACH

A non parametric method does not assume the returns of the assets to be distributed according to a specific probability distribution. It is a simple way to calculate the Value at Risk that is widely used just because the fact that it ignores many problems and still produces relatively good results (Dowd 1998). The non parametric historical simulation approach used in the study is the historical simulation approach.

The idea of the historical simulation approach is to use the historical distributions of price changes in order to calculate the VaR. Historical data is collected for a chosen period and then the historical price changes are assumed to be a good assessment of future price changes. The function for historical simulation can be described as below (Goorbergh & Vlaar 1999).

$$\widehat{VaR}_{t+1|t} = -W_0 R_t^P$$

The $\widehat{VaR}_{t+1|t}$ is the VaR value at time t+1 where W_0 is the initial value of the asset and R_t^P is the p:th percentile of each subsample, which means taking the asset return between time t and t_{n-1} preceding returns and calculating the percentile of these values that corresponds to our selected confidence level.

The merits of this approach, according to Dowd (1998), are its simplicity which makes it a great help for risk managers and the data needed should be available from public sources. Another important benefit is that it does not depend on the assumption about the distribution of returns. Whether it is the normal distribution or not that gives the best estimates is not the question here. Even though the approach is indifferent to which distribution the returns have is it important to remember that it does assume the distribution of returns to remain the same in the future as it has been in the past. Dowd (1998) states that this makes the approach less restrictive because we neither have to assume

[1] Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Sweden, the United Kingdom and the United States.

that the changes are independent giving the approach no problems in accommodating the fat tails that torment normal approaches to VaR.

A possible disadvantage, argued by Jorion (2001), is that the approach requires a lot of data to perform well at higher confidence levels. When estimating VaR at a 99% confidence level, intuitively at least 100 historical values have to be inputted. But even then the approach only produces one observation in the tail. Perhaps not enough historical data is available to produce a good VaR estimate, a problem that on the other hand can occur for most of the VaR approaches. Jorion (2001) further argues that there is a trade off in the approach when including more and more historical values. On one hand, the approach becomes more accurate at higher confidence levels, but at the same time, the risk of including old values that are not relevant for future returns is higher.

The historical simulation approach also assumes that the past is identical to the future, making historical risks the same as future risks, which might not be the case if there have been extra ordinary events in the recent past. Perhaps this is most often not the case, but the importance of getting data that truly reflects the past becomes critical. If the chosen data period is too short, it might reflect a more or less volatile period that does not give a good picture of the historical volatility. It might include periods of unusual kind that is not representative.

Goorbergh & Vlaar (1999) argues that it is not generally possible to make historical simulation VaR predictions on window sizes smaller than the reciprocal of the selected confidence level. This means that for evaluating a 99.9% confidence you would need a window size of at least $\frac{1}{1-0.999} = 1000$ observations.

The conclusions that can be drawn is that the approach does have both advantages and disadvantages which might make it important to complement it with other tests that pick up plausible risk not represented in the historical data.

2.3.2 MOVING AVERAGE APPROACH

A parametric approach assumes that the assets return follow a probability distribution. The more naïve approach, here called the moving average approach, is a parametric approach that assumes that the returns of assets are normally distributed. The approach measures the standard deviation of past returns and uses this together with the standard normal distribution to describe the probability of different outcomes in the future. However this approach as mentioned above disregards the well-established phenomenon of volatility clustering. Dowd (1998) sees this as a big problem because there is a tendency for high- volatility observations to be clustered with other high-volatility observations and the same is true for low-volatility observations that cluster with other low-volatility observations. This phenomena has been confirmed by many studies since the first observation made by Mandelbrot (1963).

Under this approach, the sample standard deviation is first calculated from a number of historical observations using this function.

$$\sigma = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(u_i - \bar{u})^2}$$

The number of observations is represented by n, u_i stands for the daily log price for day *i* and $\bar{u}$ is the average daily log price.

Calculating the actual VaR value is then done by retrieving the number of standard deviations that corresponds to the selected confidence level from the normal cumulative distribution function presented below, and multiplying it with the standard deviation of the historical observations (Jorion 2001).

The function below is the cumulative distribution function (cdf) of the standard normal distribution. This function calculates the probability that a random number drawn from the standard normal distribution is less than *x*. The inverse of this function gives the value of x that is the limit under which a random value drawn from the standard normal distribution will fall with the inputted probability (Lee, Lee & Lee 2000). Erf stands for the error function.

$$cdf = \frac{1}{2} \times \left(1 + erf\left(\frac{x}{\sqrt{2}}\right)\right)$$

Based on this the confidence levels that we have chosen corresponds to the following number of standard deviations (Penza & Banzal 2001).

Confidence level	Standard deviations
95,00%	1,645
99,00%	2,326
99,90%	3,090

Table 2.3.2 Relating standard deviation to chosen confidence level.

An important part of this approach is choosing the number of days to base the volatility estimate on. If the volatility of the underlying asset is estimated from too many historical observations, there is a risk of including observations that are too old and irrelevant. On the other hand, a historical value based on too few vales means a risk of getting values that are too heavily influenced by recent events. As a rule of thumb, Hull (2008) mentions that you should base your historical volatility estimate on equally many days as you are trying to forecast into the future, but no less than 30. In this study, the number of days would therefore be 30.

The distribution of the returns can be different for varying assets. The most widely used distribution is the *normal distribution* because of its simplicity, however, it does seldom exactly fit the returns of an asset. Dowd (1998) states that normality gives us a simple and tractable expression for the confidence interval for the VaR estimate, and with a normal return, the VaR is just the multiple of the standard deviation. It is only about two parameters, the holding time and the confidence level. These parameters can vary depending on the user. Also, going from a VaR estimate for a specific confidence level to another one can easily be done just by changing the numbers from the normal distribution associated with the confidence level.

2.3.3 GARCH APPROACH

GARCH is not really a VaR method in itself but rather a more advanced way to compute the standard deviation of past returns that in turn is intended to give more precise VaR estimates. The GARCH approach is considered to be semi parametric, meaning that it has both parametric and none parametric properties. It is parametric in the sense that the estimated variance of the approach is used with the normal distribution, but it is at the same non parametric since the inputted variables are real returns and not estimated volatilities.

The GARCH approach for variance, h, looks like this (Jorion 2001):

$$h_t = \omega + \alpha r_{t-1}^2 + \beta h_{t-1}$$

Where ω, α and β are weights. α is the weight for how much of today's variance influences our forecast, β weights how much of yesterdays estimation weights into the forecast. Jorion (2001) points out that the sum of α and β must be less than one and is usually called the persistence of the approach. It is called persistence, since during a multi period forecast into the future, the weights would decay as they are both less than one and the calculated value of variance would revert back towards the estimated long term variance ω.

Though the use of its non parametric input, the GARCH approach can takes volatility clustering into consideration. This means that it takes into account that large changes in prices at a specific time are usually followed by more large changes and vice versa. This is done by weighting historical observations and recent observations unequally.

There is no way of solving for the correct values of ω, α and β mathematically so it has to be done numerically using maximum likelihood estimation (MLE). To estimate the parameters ω, α and β the logarithm of the likelihood function of the normal distribution is used (Jorion 2001).

$$\max F(h_t) = \sum_{t=1}^{T} \left(ln\left(\frac{1}{\sqrt{2\pi h_t}}\right) - \frac{r_t^2}{2h_t} \right)$$

ω, α and β gives the values of h_t, so maximizing this function by adjusting ω, α and β gives us the optimal values of ω, α and β. This optimization is done for each and every one of the historical observation on which the VaR measure is to be based. The parameters ω, α and β are then adjusted to give a maximum value of all of those observations combined for the chosen window in time.

It should be noted that as the MLE function is derived from the normal distribution function, it also assumes that the returns of the assets are normally distributed. The MLE function may therefore, if applied to none normally distributed returns, produce questionable results (Jorion 2001).

2.4 THE UNDERLYING ASSETS

The underlying assets, on which our VaR calculations are based, are chosen because of their fundamentally differences which would indicate fundamentally varying return distributions. All three assets show different properties making them interesting for comparison according to the purpose of the study.

Whenever holding an asset you face the risk of the asset gaining or losing value in relation to its purchase price. The field of risk management is all about assessing and mitigating risk. An effective risk management is according to Saunders & Cornett (2007) central to a financial institutions performance. This is also true for any company exposed to risk. As discussed in the problem discussion VaR is a useful measure for all of our assets.

2.4.1 BRENT OIL

Crude oil is the most traded commodity on the world market and the most important hubs for this trading are New York, London and Singapore. The prices of the Brent oil, which is the world benchmark for crude oil and is pumped from the North Sea oil wells is the oil used in the study. According to Eydeland & Wolyniec (2003), about two thirds of the world supply of oil is priced with references to this benchmark.

The oil price is very volatile, and a reason for this is the influence of *Organization of the Petroleum Exporting Countries* (OPEC). This is an oil cartel that was created in 1960 and now consists of twelve member countries. Their oil production stands for nearly half of the world production, and by raising or lowering their production, they can affect the price. There are different opinions about OPECs influence on the balance of demand and supply. Some people claim that they prevent the forces of the market to set the real price while others say that they can only affect the supply of oil, not the demand, making the forces of the market through the demand create an equilibrium price.

2.4.2 OMXs30

OMXs30 is an abbreviation of OMX Stockholm 30 and is the denomination of the 30 most traded stocks on the Stockholm stock exchange. It is a capital weighted index that measures the price development of the stocks. The value of the stock for the separate companies is the basis of their part of the index. The assembly of the index is conducted twice per year and the revenues of the companies then work as a ground for the companies elected. (www.omxnordicexchange.com)

2.4.3 THREE MONTH SWEDISH TREASURY BILLS

The Swedish treasury bills are issued through the Swedish National Debt Office and the bids are then submitted to dealers authorized by the same. The bills are traded on the secondary market, which is considered to be a healthy market when looking at the liquidity. The maturity varies within a year, however, the treasury bills used in this study are the three month bills (STB3M). (www.riksbank.se)

2.5 BACKTESTING WITH KUPIEC

When calculating the VaR, one is only interested in the left tail that represents the cases when the returns are worse than expected. To evaluate the measures used, a back- test can be conducted. Goorbergh & Vlaar (1999) explains the test as counting the number of days in the evaluation sample that had a result worse than the calculated VaR. An observation where the actual return exceeds the VaR is called a VaR break. The back test is important to use, and by performing it and comparing the different approaches, one can ensure that the approaches are properly formed according to Costello & Asem & Gardner (2008).

Jorion (2001) state that the number of VaR breaks is expected to be the same as one minus the level of confidence. So for a sample of 100 observations where a 95% confidence VaR is calculated, we would expect five $(100\% - 95\% \times 100 = 5)$ VaR breaks to occur. In the study this is called the target number of VaR breaks. If there are more or less VaR breaks than expected, it is because of deficiencies in the VaR approach or the use of an inappropriate VaR approach. A widely used back test is the Kupiec test. This test uses the binominal distribution to calculate the probability that a certain number of VaR breaks will occur given a certain confidence level and sample size. The Kupiec test function is (Veiga & McAleer, 2008):

$$P(x|n,p) = \binom{n}{x} p^x (1-p)^{n-x}$$

The variable x is the number of VaR breaks, N the sample size and p corresponds to the level of confidence chosen for the VaR approach (making a 95% confidence level a 5% probability input for the approach). If the sample size is inputted and p is set to one minus the level of confidence, the binomial function produces the likelihood that a specific number of VaR breaks is to occur. By using the cumulative binomial distribution, it is possible to calculate an interval within which the number of VaR breaks must fall, in order for the approach to be accepted. This is done by calculating for which values of N that the cumulative binomial distribution produces probabilities that are in the interval $2.5\% < P < 97.5\%$ (which corresponds to a 95% test confidence). VaR approaches that produce values of N that lies within this range can therefore be accepted. If the approach produces values of N outside this span, the approach is rejected. A rejection means that the confidence level that we used in the VaR approach did not match the actual probability of a VaR break. This in turn indicates that the approach is not performing well and that it should be rejected.

3. METHOD

In this chapter, the methods used in the study are presented. First, the nature of this study is described and then the calculations of VaR are explained for the different approaches.

3.1 ANALYTICAL APPROACH

This study takes an analytical approach and assumes an independent reality. This is a method approach that is widely used and the history of it goes way back in time. The most distinguishing characteristic concerning the procedure of this approach is, in accordance with Arbnor & Bjerkne (1994), its cyclic nature. It starts with facts, ends with facts and these facts can be the start of a new cycle. The meaning is to extract good models in order to describe the objective reality. The models within this approach are most often of quantitative character and this is also true for this study. It is natural that logic and mathematics has a part to play in the search for the best model to apply.

3.1.1 QUANTITATIVE APPROACH

In this study, the tests are based upon a lot of historical data. Daily price changes are studied for three different assets that are widely traded, this form of study makes us apply a quantitative approach. This approach compared to a qualitative approach is according to Arbnor & Bjerkne (1994) often more about clear variables and can cover a larger number of observations, which fits our composition better. It further assumes that we can make the theoretical concepts measurable. In the past, the quantitative approach has in many cases wrongly been seen as the only real scientific method, or as Holsti said in 1969, "If you can't count it, it doesn't count." There are however limitations connected to the approach and Holme & Solvang (1991) points out that if you are not aware of them, the result can easily be misjudged. Numbers are often seen as the truth, and this trust can create problems making the analysis of the numbers very important.

Historical data have been collected in this study, and based on these, a number of calculations have been made in order to extract the VaR. The purpose is to be able to see which approach that performs best, and the result is based on the figures from the calculations. However the result is not just a number it is also put into context to better get an understanding of its meaning.

3.1.2 DEDUCTIVE APPROACH

When taking a deductive approach one starts from an already existing theory and then either strengthen or weaken the confidence of the theory, meaning that one starts from a general law and then moves to a separate case. In our study, we use three well known approaches for calculating VaR and test their performance on assets with different return characteristics. The purpose is to test models, not to create new ones.

As a result of this study, the function of the approaches will perhaps be strengthened because of a specific character, however, the same approach applied on the same asset can turn out to be unreliable for a certain confidence level. The approachs can from the start be rather simple, however, by adding a dimension like the GARCH approach, one can make the approach more complex and maybe better fitted to a certain asset. The test of the approaches will be based on backtesting, we will count the numbers of VaR breaks, meaning the number of times the loss exceeded the calculated VaR. Reasons for the actual number of VaR breaks will be analyzed.

3.1.3 RELIABILITY

In order to decide the reliability of the results, one should be able to do the test a number of times to see if the same results are achieved. Arbnor & Bjerkne (1994) calls this a test-retest, and it is helpful in telling whether the study and its results are reliable. All the data used in the study to calculate the VaR is public information, making the test reproducible.

3.1.4 VALIDITY

The most important factor when judging whether the results are justifiable in the analytical approach is through the validity. If you cannot answer the question whether the results gives a true picture of the reality, the results can easily become meaningless. The closer we get to the true situation, given a specified definition, the higher the validity. The relation between theory and data is vital, and Holme & Solvang (1991) points out that the validity can be enhanced by a continuously adaption between the theories and the methods used in the examination. This is done by choosing methods that has the ability to handle or show the impacts that different assets characteristics can have on the measured VaR. Arbnor & Bjerkne (1994) stresses the importance of three kinds of validity and the need for combining them. These are the surface validity that deals with the reasonability assessment in relation to earlier results, the internal validity dealing with the question if we had expected the result we have concluded and finally the external validity and thereby the usefulness of the result in other areas. These questions will be further dealt with in the analysis.

It is further important to have in mind that an approach of the kind used in the study can give you a number of the VaR but it is always an estimation of a possible future loss. Given a certain confidence level the VaR can be calculated, however, the number received is not the absolute truth, it cannot be guaranteed that the future losses won't be larger than expected.

3.2 CALCULATION OF VaR

With many different approaches and models, the choice that VaR users face is the choice of picking the right one that matches their purpose best. The approaches should make estimates that fit the future distribution of returns. If an overestimation of VaR is made, then operators ends up with an overestimate of the risk. This could result in the holding of excessive amounts of cash to cover losses, as in the case with banks under the Basel II accord. The same goes for the opposite event, when VaR has been underestimated resulting in failure to cover incurred losses.

In this essay, we will be calculating VaR for three different underlying assets and compare the results. The assets that will be used for our calculations are Brent Oil, OMXs30 stock index, and three month Swedish treasury bills. The VaR approaches that we will calculate are the historical simulation, the moving average approach, assuming normal distribution, and the GARCH approach, assuming normal distribution. When using the parametric approaches, we do suspect that the returns that we have based the calculations on are most likely not perfectly normally distributed. In fact, economic time series rarely are (Jorion 2001). The performance of these parametric approaches will therefore partly be determined by how well the normal distribution assumption fits the actual distribution of the returns.

The tool that we have used to perform our calculations is Microsoft Excel. Excel is a very versatile tool that can be helpful if one know how to use it right. However, the downside of Excel is that it is much slower than software that is specifically designed for financial calculations. There are also no automated functions for calculating a lot of the more advanced operations, often used in financial time series analysis, such as autocorrelations etc. These more advanced calculations have to be done manually, which can be very time consuming and also limits us somewhat when it comes to computing more advanced approaches.

3.2.1 HISTORICAL SIMULATION APPROACH

Calculating a VaR measure using the historical simulation approach is not mathematically complex but can be trying since the calculations require a lot of historical data. The first task is to pick a historical time frame on which to base the simulation. The more historical values we base our calculations on the more observations we will get in the tails and thus the more accurate the VaR measure will be at higher confidence levels.

The downside is that adding more historical data means adding older historical data which could be irrelevant to the future development of the underlying asset. Our simulation will be based on a sliding window of the previous 2000 observations which corresponds to 8 calendar years. We have selected this rather large window for the historical simulation approach as we want the approach to perform well at the 99% and 99.9% confidence levels. As argued in the theory chapter it makes no sense to set the window size to any less than 1000 observations for the 99.9% confidence level, but even at this level we theoretically only have one observation in the tail. We therefore decided to set the window size to 2000 observations to enhance the performance of the approach at higher confidence intervals.

Extracting the VaR measure from the historical data simply requires us to choose the desired confidence level and pick out the n:th observation in the historical data that corresponds to that confidence level. For example, a 95% confidence level means that we are interested in the worst 5% of the observations. If we for example are using 1000 observations, this would mean that the 95% VaR would be the 50th worst observation ($1000 \times 5\% = 50$).

The PERCENTILE –function we used in Excel calculates the n:th percentile of the values in a chosen data set. If the desired percentile is not a multiple of the number of values in the dataset, Excel will do a linear interpolation between the two closest values to find the desired value.

3.2.2 MOVING AVERAGE APPROACH

Under this approach, we continuously measured the standard deviation of returns over a window of the last 30 days. This standard deviation was then multiplied with the number of standard deviations extracted from the standard normal distribution that corresponds to the selected confidence level. In other words, the standard deviation changes as the window moves along but the mean is assumed to be zero.

3.2.3 GARCH APPROACH

The challenge with the GARCH approach is to estimate the parameters ω, α and β. As described in the theory chapter, these parameters has to be solved for numerically, using maximum likelihood estimation. The literature that we have studied describes the MLE function but provides little practical information on how to implement it. The math behind maximum likelihood estimation can be complicated to understand for those not familiar with business statistics. Many analysts that do these types of calculations use preconfigured software and seldom have to engage in the mathematics that lies behind the calculations.

The MLE formula itself is not complicated and was implemented for each and every observation in our study. Maximizing the function was done using the SOLVER function in Excel. The first problem that we encountered was how to decide how many historical observations that we should base our MLE on. At first, we had planned to base the MLE on a moving window of 250 values, matching a year in trading days, but this approach produced very inconsistent results often setting one of the parameters, α or β, close to zero and the other close to one. The conclusion we drew was that the MLE was based on too few values to be consistent. As mentioned in the theory chapter, the MLE function also assumes that the returns are normally distributed. Thus, the smaller the sample that the MLE estimation is based on, the larger the risk that those values might be significantly parted from the normal distribution.

We looked through a great deal of articles and books in hope of finding some kind of documentation on how many observations to include in our MLE but we were unable to find any recommendations to go by. We therefore made our own estimation of an appropriate time frame on a trial and error basis. This was done by constructing a macro in Excel that would maximize the MLE function, first based on the 250 first observations and then add 100 observations at a time each time re-estimating the parameters ω, α and β until we had tested all of the observations. The macro is shown in the appendix 1. The resulting values of α, β and ω are shown in the graphs 3.2.3a, b and c.

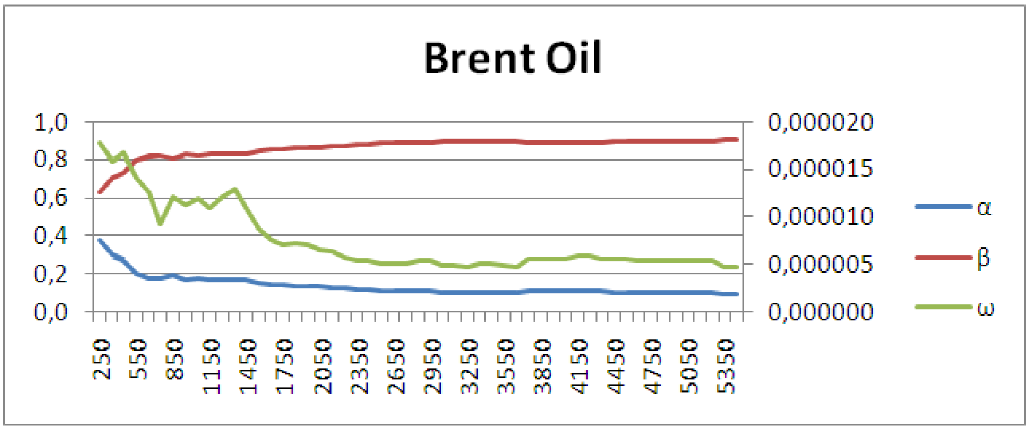

Figure 3.2.3a MLE estimation graph for Brent Oil

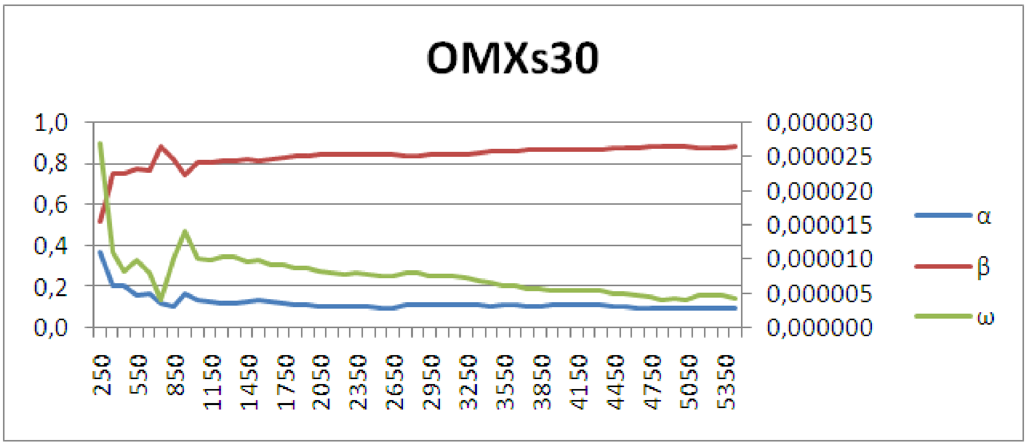

Figure 3.2.3b MLE estimation graph for OMXs30.

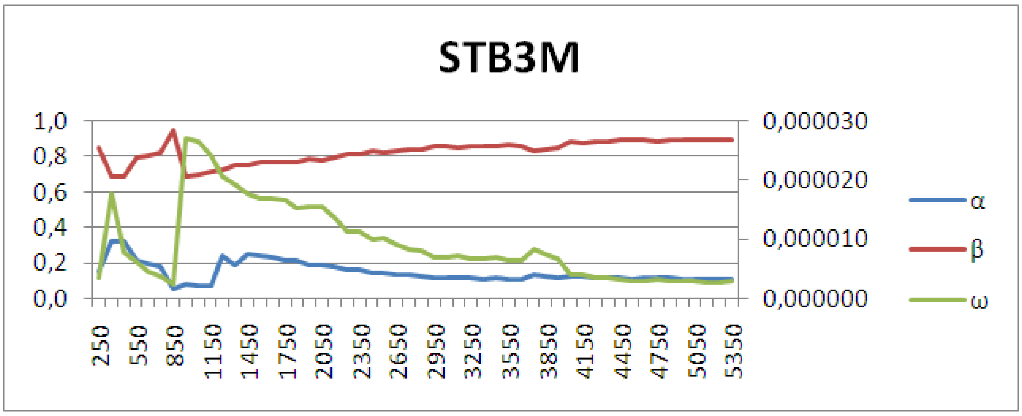

Figure 3.2.3c MLE estimation graph for STB3M.

The graphs show the values of α and β in blue and red on the left vertical axis and the values of ω in green on the right vertical axis. The horizontal axis shows the number of observations included in the estimate. Since the sum of α and β must be less than one and therefore α must go down for β to go up. This is why the values diverge during the first estimations to later level out or become stable.

One can clearly see how values of ω, α and β stabilize as we add more and more values to the MLE. We also interpret these results as an indication of to what degree the returns are normally distributed or not.

By looking at the graphs, we see that at around 2000 observations the MLE has stabilized for all of our three assets, which lead us to conclude that 2000 observations is a suitable number of observations to base our MLE on.

After deciding the appropriate number of observations to base our MLE on, we then re calculated the MLE function for our datasets, now with a lag of 2000 observations from 1996-01-01 onwards. New graphs have been plotted for the values of ω, α and β, and they can be found in appendix 3.

After the values of ω, α and β had been estimated, they were inputted into the approach and the results will be shown in the results and analysis chapter.

3.3 SKEWNESS

The skewness tells us whether the data is symmetric or not. Normally distributed data is assumed to be symmetrically distributed around its mean, i.e. it has a skew of 0. A dataset with either a positive or negative skew therefore deviates from the normal distribution assumptions. This can cause

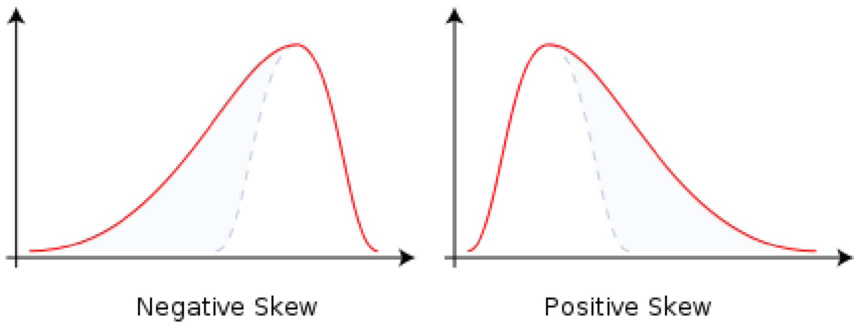

Figure 3.3 Graph showing a negative respectively a positive skew.

parametric approaches, in our case the moving average approach and the GARCH approach, to be less effective if asset returns are heavily skewed, since these approaches assume that the returns are normally distributed. The result can be an overestimation or underestimation of the VaR value depending on the skew of the underlying assets return (Lee, Lee & Lee 2000).

3.4 KURTOSIS

The kurtosis measures the *peakedness* of a data sample and describes how concentrated the returns are around their mean. A high value of kurtosis means that more of the data's variance comes from extreme deviations. The kurtosis of the normal distributions also mentioned by Lee, Lee & Lee (2000)

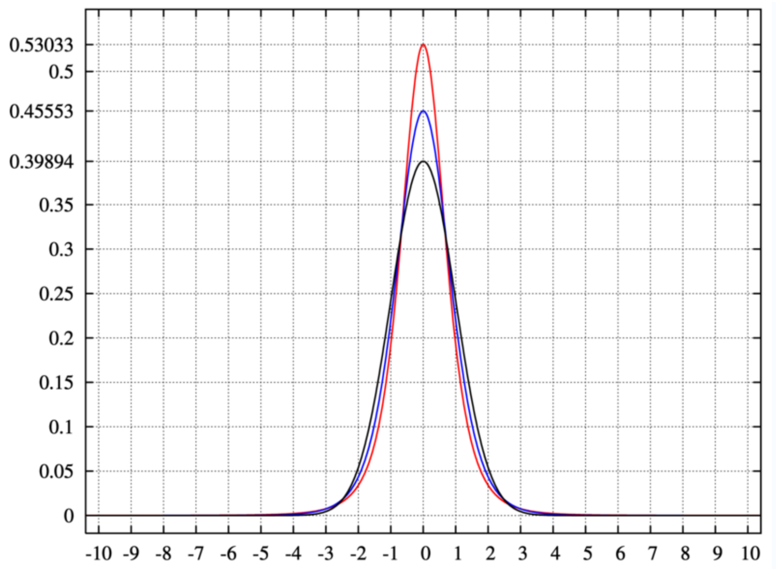

Figure 3.4 Probability density function for the Pearson type VII distribution with kurtosis of infinity (red); 2 (blue); and 0 (black).

as mesokurtic is three so as with the skewness any deviations from the normal distribution cause problems for the parametric approaches. A high kurtosis means that the assets returns consist of more extreme values than modeled by the normal distribution. This positive excess kurtosis is according to Lee & Lee & Lee (2000) called leptokurtic and lower kurtosis meaning a negative excess is called platykurtic. Low kurtosis produces VaR values that are too small.

3.5 THE SOURCE OF DATA

The data used in the study are time series that reflects the historical price changes for the chosen assets. The period, which the calculations are based on, stretches from 1987-01-01 to 2008-09-30 for all the three assets. However, the data from 1987 to 1995-12-31 has been collected to enable us to base our calculations that require historical data on.

Different sources have been used to collect the data. Daily oil prices are taken from the Energy Information Administration (www.eia.doe.gov), who provides official energy statistics from the US government. The OMXs30 data comes from OMX Nordic Exchange (www.omxnordicexchange.com), and the Treasury bill data can be found on the homepage of the Swedish central bank (www.riksbank.se). The characteristics of the different assets are described below.

We have calculated some characteristic key figures about the data sets in order to better be able us to compare them. The key figures that we have calculated are skewness, kurtosis, volatility, and average price change. Table 3.5 shows some of the statistical characteristics of each dataset, measured during the entire period of 1987 – 2008.

	Brent Oil	OMXs30	3MSTB
Skewness:	-0.83	-0.02	3.77
Kurtosis:	19.14	7.28	182.51
Daily volatility:	2.32%	1.47%	1.49%
Annual volatility:	36.78%	23.37%	23.66%
Average log price change:	1.63%	1.03%	0.58%

Table 3.5 Statistical characteristics of the asset returns.

3.6 BRENT OIL

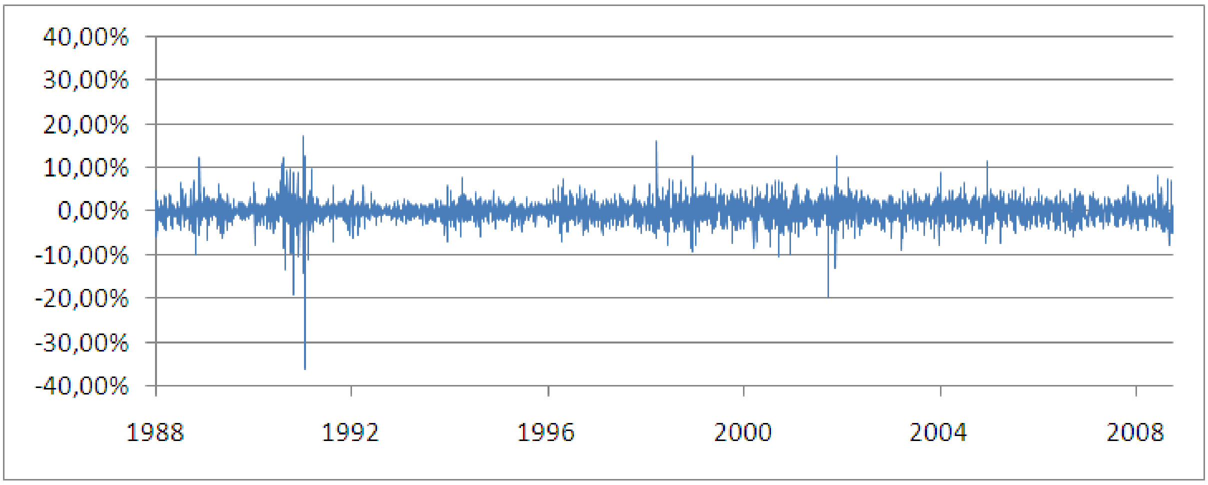

Figure 3.6a Daily returns for Brent oil.

The oil prices are the most volatile prices in this study, with a 2.32% daily volatility and 36.78% annual volatility. The kurtosis is 19.14, which make it far narrower with much fatter tails than the normal distribution.

The skewness of the Brent oil returns have a negative skew of -0.83, which indicates that the returns are skewed to the right. This means yet another departure from the assumptions of normality in the parametric approaches causing these to perform less effective. The daily average log price change is 1.63%, which basically means that over the measured time period the average return was 1.63%.

In graph 3.6b, we have fitted a normal distribution probability distribution curve against a histogram of the returns. The most visually prominent deviation from the normal distribution assumption is the kurtosis, which can be seen from the middle bars of the histogram rising above the normal distribution. There are also several outliers to right and left of the histogram (which can be hard to see because of the scale of the picture) which lay outside the bounds of the bounds of the normal distribution curve.

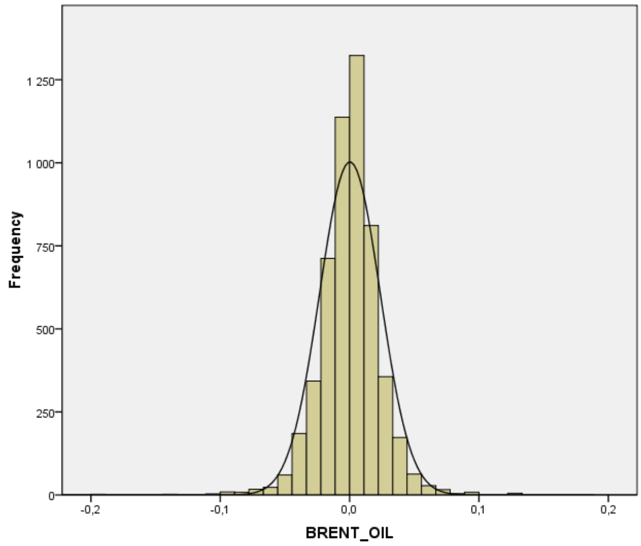

Figure 3.6b Histogram showing the daily returns combined with a normal distribution curve.

3.7 OMXs30

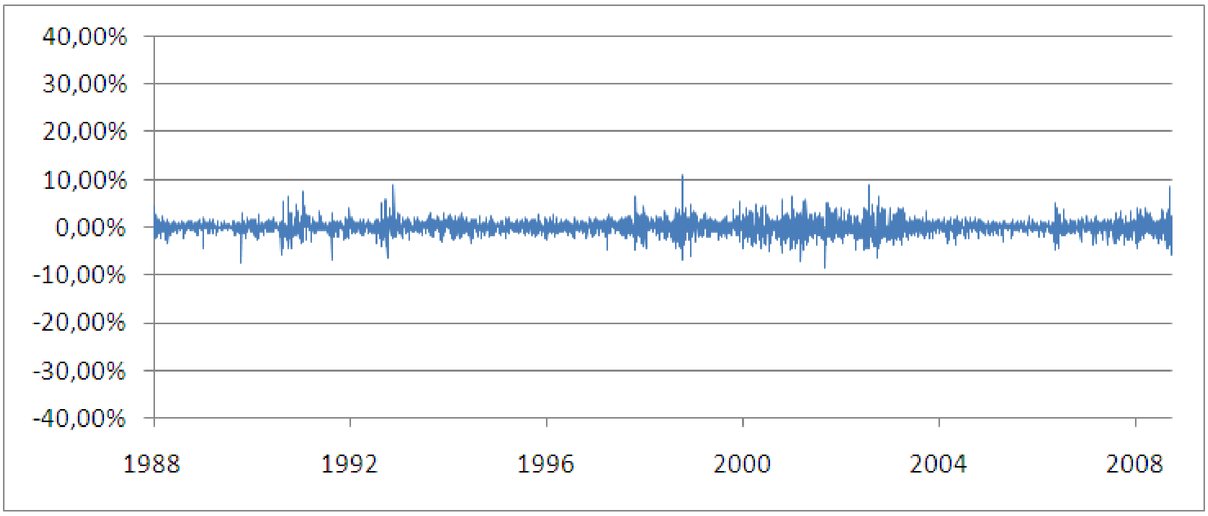

Figure 3.7a Daily returns for OMXs30.

Since the returns of the index corresponds to a well diversified portfolio, they are not very volatile. The average volatility is 1.47% daily and 23.37% annually. The kurtosis is 7.28, which suggest that the distribution of the values is narrower than the normal distribution and with somewhat fatter tails. It is however not as high as it was for the Brent oil, which might indicate a better compatibility with the parametric approaches, at least at lower levels of confidence.

The skewness is only -0.02, which and the average daily log price change is 1.03%. Apart from the kurtosis the normal distribution curve is a pretty nice fit on these returns.

There are however large outliers in that are not clearly visible in the graph which lay outside the confinements of the normal distribution curve. There is also no significant sign of autocorrelation as shown by the table in the previous chapter.

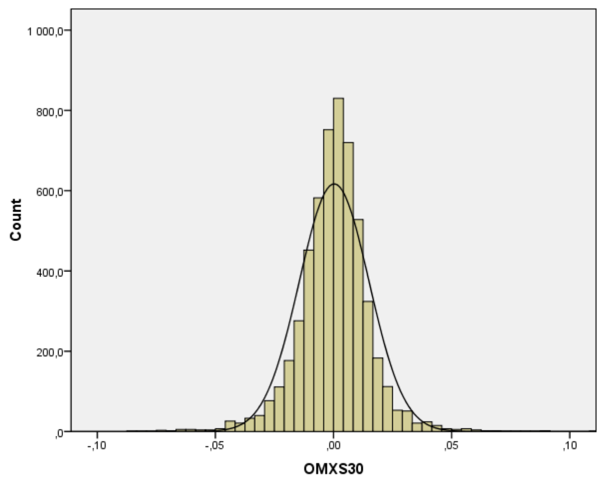

Figure 3.7b Histogram showing the daily returns combined with a normal distribution curve.

3.8 THREE MONTH SWEDISH TREASURY BILLS

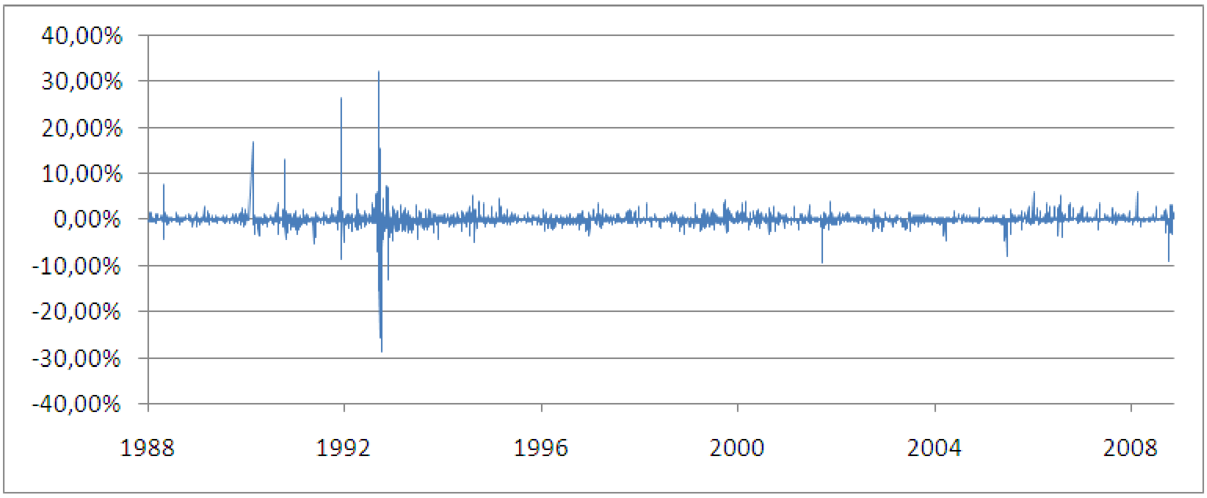

Figure 3.8a Daily returns for STB3M.

The returns of the treasury bills have the most extreme characteristics of all the datasets. It has a strong positive skewness of 3.77, and a very large kurtosis of 182.51 suggesting a poor fit to the normal distribution which also can be observed from the histogram in figure 3.8b . The average log price change is 0.58% but from figure 3.8a we can see that volatility varies from low to extreme. The most extreme returns are from the early 1990's when the banking crisis occurred in Sweden.

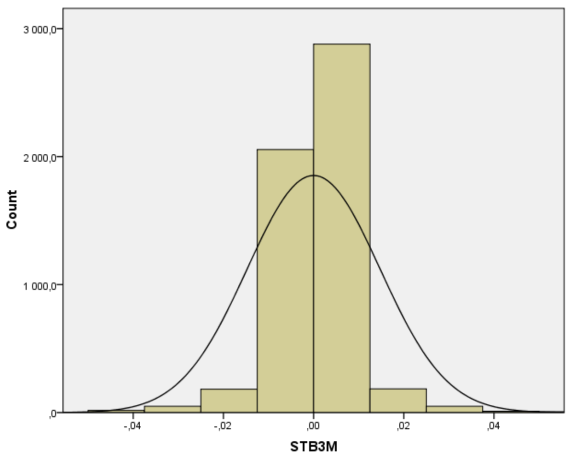

Figure 3.8b Histogram showing the daily returns combined with a normal distribution curve.

3.9 AUTOCORRELATION

In time series analysis, autocorrelation, or serial correlation as it is also called, is a measurement of whether the data correlates with itself over time. Autocorrelation in financial time series data is measured by regression of the observed returns with a lagged version of themselves. Regressing the returns with themselves like this can be described as testing if it is possible to describe returns of today as a linear function of the returns from yesterday. The resulting residuals from the regression are then tested for autocorrelation. The presence of autocorrelation means that the applied approach will have a poor fit to the actual data which leads the analyst to conclude that the returns of today cannot be accurately described as a linear function of the returns of yesterday. Thus there must be other factors besides the historical returns that effect today's returns and thereby approaches that try to predict future returns by looking at the past will be less effective (Tsay, 2002).

We have chosen to perform a one lagged Durbin Watson (DW) test on our data to see if there are any first order autocorrelation. The DW test uses the following formula to calculate its statistic:

$$DW = \frac{\sum_{t=2}^{T}(e_t - e_{t-1})^2}{\sum_{t=1}^{T} e_t^2}$$

The parameter e denotes the residuals from the regression. The DW statistic can assume a value between 0 and 4. Values of the DW statistic larger than 2 indicate negative autocorrelation whereas values less than two indicate positive autocorrelation. A value close to 2 indicates that the residuals are probably not auto correlated (Tsay 2002).

When testing for autocorrelation, a confidence interval is calculated around 2 within which the DW statistic must fall in order for the residuals to be considered not to be not correlated with the given degree of significance. Harvey (1990) has shown that for large sample sizes, this interval is approximately normally distributed with a mean of two and a variance of $\frac{4}{N}$.

We have used the statistical software program R to test our time series for autocorrelation. The autocorrelation tests results have been compiled in table 3.9 and some additional regression statistics can be found in appendix 2.

RESULTS FROM AUTOCORRELATION TEST							
Asset	No.	Lag	Autocorrelation	dL	D-W	dU	P
BRENT OIL	5 454	1	0,000660287	1,955455	1,998650	2,044545	96,40%
OMXs30	5 475	1	0,000758127	1,955540	1,997565	2,044460	94,80%
STB3M	5 468	1	0,009438294	1,955512	1,980684	2,044488	36,80%

Table 3.9 Results from autocorrelation test.

The p value is a probability ratio that R calculates which donates the probability that the DW-statistic would occur randomly. A low p value thus indicates that the DW statistic in question is very unlikely if the approach is correct. The null hypothesis is:

$$H_0 = Returns\ are\ not\ autocorrelated$$
$$H_a = Returns\ are\ autocorrelated$$

The null hypothesis will be rejected if the p value is less than one minus the selected confidence interval which in our case is 95%. Thus if p would have been less than 5%, the null hypothesis would have been rejected, which would have meant that the returns were autocorrelated. As stated by the p values and the DW statistics in the table, there is no sign of autocorrelation in the residuals.

3.10 BACKTESTING WITH KUPIEC

In order to measure which of the VaR approaches that has been most successful, we will back test the VaR values against the actual returns of each day. In the event that the actual loss on a particular day is less than the projected VaR value, we say that a VaR break has occurred.

The expected number of VaR breaks is one minus the selected level of confidence times the number of observations. The numbers have been rounded off to the nearest integer. We have compiled a table of the desired number of VaR breaks in table 3.10 .

DESIRED NO. OF VaR BREAKS AND KUPIEC TEST INTERVALS										
		95%			99%			99.9%		
Asset	No.	MIN	TARGET	MAX	MIN	TARGET	MAX	MIN	TARGET	MAX
Brent Oil	3 259	139	163	187	22	33	43	0	3	6
OMXs30	3 215	137	161	184	22	32	43	0	3	6
STB3M	3 225	137	161	185	22	32	43	0	3	6

Table 3.10 Target number of VaR breaks and Kupiec test intervals.

The nearer the desired no of VaR breaks an approach produces the better the approach is considered to perform. Whether the produced number of VaR breaks is greater or smaller that the target value makes no difference, it is the absolute number of VaR breaks that the approach deviates with that determines how effective it is.

4. RESULTS & ANALYSIS

In this chapter, we will conclude our results and thereafter analyze them in relation to the purpose of the study. The analysis will be divided, treating the approaches separately, in order to get a better overview of the approaches and their results.

4.1 BACKTESTING RESULTS

The backtesting results are presented for the three different assets under the three approaches. The outcomes depend on the characteristics of the asset, and by examining them, the result can be explained.

4.1.1 HISTORICAL SIMULATION APPROACH

The backtesting results for the three assets, calculated using the historical simulation approach, are shown in table 4.1.1a. The table shows the minimum and maximum values allowed by the Kupiec test, the target number of VaR breaks and the resulting number of VaR breaks from the simulation. VaR breaks that are within the range allowed by the Kupiec test are marked green and those that are not are marked red.

KUPIEC TEST - HISTORICAL SIMULATION, 2000 DAY LAG													
		95%				99%				99.9%			
Asset	No.	MIN	TARGET	RESULT	MAX	MIN	TARGET	RESULT	MAX	MIN	TARGET	RESULT	MAX
Brent Oil	3 259	139	163	209	187	22	33	37	43	0	3	7	6
OMXs30	3 215	137	161	216	184	22	32	54	43	0	3	5	6
STB3M	3 225	137	161	109	185	22	32	24	43	0	3	7	6

Table 4.1.1a Backtesting results with historical simulation approach.

The results show that the approach performs poorly on the 95% confidence level while it produces better results on the higher confidence levels. On the 95% confidence level, the approach produces too many VaR breaks, which indicate that the approach overestimates the VaR On the 99% level this approach performs equally compared to the GARCH approach and outperforms the moving average approach. When it comes to the highest level this approach is the only one that comes even close to being accepted by the Kupiec test.

This approach looks at the latest 2000 daily returns when calculating the volatility, which in turn makes the approach react slowly to new information and changes in the daily returns. There is a tradeoff between new information and old historical information that we presented in earlier chapters. If a shorter interval would have been chosen, this would have increased the impact of new observations, putting higher focus on recent market conditions.

Table 4.1.1a shows the VaR estimated by historical simulation at the 99.9% confidence level for Brent Oil. Out of all the assets, the historical simulation approach performs best results with Brent oil. It does however seem that the approach generally underestimates the Value at Risk, thus producing too many VaR breaks. The characteristics of Brent oil showed that the normal distribution, and therefore the parametric approaches, would make a poor fit and that the historical simulation would suit the asset better. Brent oil is the most volatile of the assets but the volatility is rather predictable, or more correctly, high but stabile. The reason to why historical simulation performs better with Brent oil is because of its stability. As discussed earlier the historical simulation approach does not assume the returns to follow a specific probability distribution, but it does assume that the return distribution is the same in the past, in our case 2000 days, as it will be tomorrow. The stability of the Brent oil returns is therefore the reason for the good results.

When testing the accuracy of an approach it is equally important with underestimations as it is with overestimations. The results for the two other assets are about the same with the difference lying in overestimations of the Value at Risk for the STB3M and underestimations of the value for OMXs30. The volatility for these assets is not as high as for Brent oil but it changes more dramatically over time. If a shorter time period had been chosen, the results for STB3M would have looked very different ignoring the great volatility changes that appeared around 1993. The historical simulation would in that case have produced better results since the fluctuations in returns would have been less.

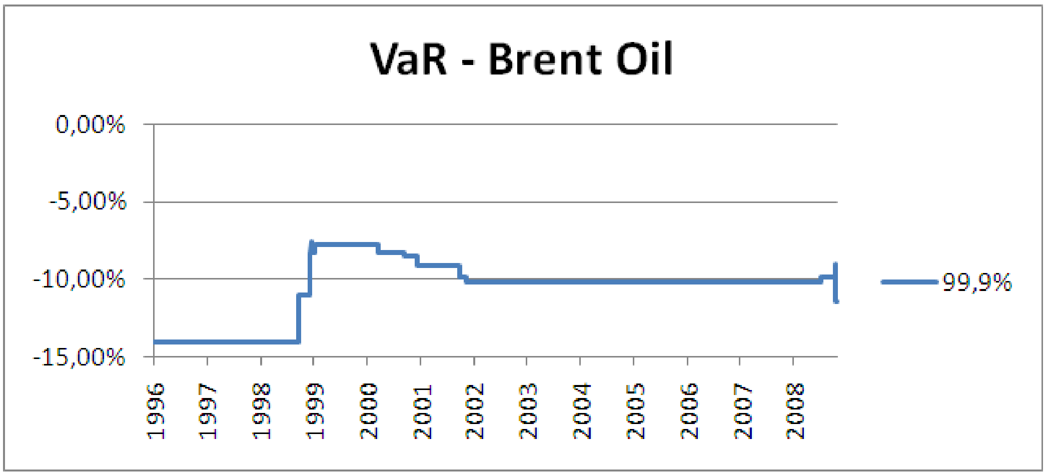

Figure 4.1.1b VaR for Brent Oil, Historical Simulation 99.9% confidence 1996 -2008.

The look of the graph 4.1.1b is characteristic for historical simulation VaR estimates. The sudden jumps up and down and the plateaus in-between signal the entrance and exit of an extreme historic return that affects the VaR as long as it remains in the window. The long window that we have chosen seems to be the cause to why the approach overestimates the VaR for STB3M and underestimates it for Brent Oil and OMXs30. At the 95%, confidence level the large window chosen seems to be resulting in too few VaR breaks for STB3M. It would seem that the large window is too large for the 95% confidence level, adding too many extreme observations in the tails. This causes the approach to overestimate the VaR, and produce too few VaR breaks. For Brent oil and OMXs30, the opposite seams to apply. The large window puts too much weight on the more frequently occurring smaller returns and draws the center of gravity in the approach away from the interesting outliers, causing the approach to produce to low VaR and thus to many VaR breaks. The same seems to apply at the 99% confidence level, but the effects seem milder, perhaps because the window size is more appropriate here.

As can be seen in table 4.4.1a the approach is rejected by the Kupiec test at the 95% confidence level. We believe this is due to the window size of 2000 observations being too large for the historical simulation approach at the 95% confidence level. The approach could probably have performed better had we chosen a more appropriate window size.

The approach does, however, perform better at the two higher confidence levels, where the selected window size is better suited. Since the historical simulation approach is nonparametric, it takes all of the outliers of the returns into account that are ignored by the parametric approaches, using the normal distribution. This is also one of the reasons to why the approach performs better at these higher confidence levels compared to the other two approaches as we will discover later.

By weighing all the returns equally, it takes time for the approach to react to fluctuations in the returns. This means that old extreme outliers in the returns affect the calculated VaR for a long time. When dealing with returns that suffer from extreme leptokurtosis as is the case with the STB3M, one ends up with an average VaR that is considerably higher than the average return, in other words the few extreme outliers have a too big impact on the VaR value given a certain window size and level of confidence. The historic simulation approach is generally considered to perform well with returns that are leptokurtic, i.e. has fat tails, but there seems to be a limit when the leptokurtosis becomes too much also for the historical simulation approach to handle. We therefore think that the size of the historical window has to be chosen with respect not only to the confidence level but also to the kurtosis.

The historical simulation approach assumes that the distribution of returns does not change over time. Without this assumption, there would be no reason at all to look at the past returns in hope of predicting the future. In other words, one contributing cause of error in the approach could be that the distribution of returns is not stationary.

Over all this approach performs best on high confidence levels, and this is because the chosen historical window size is better suited for these confidence levels. The reason that the approach performs relatively better at high confidence intervals compared to the other three approaches is that it considers the extreme values that falls out of the normal distribution. This can clearly be seen in the histograms in chapter three shoving the daily returns and the normal distribution. For Brent oil and the OMXs30, a significant number of observations can be found in the tails. For the STB3M, most observations fall in the middle of the normal distribution, causing the approach to overestimate the VaR. Historical simulation can be recommended when estimating VaR at high levels of confidence for returns that are stationary and not too leptokurtic.

4.1.2 MOVING AVERAGE APPROACH

The moving average approach is the most basic approach for calculating VaR and it has been around the longest. Even though it is elementary, it performs quite well at lower confidence levels and under the right conditions. The back test results are shown in table 4.1.2 .

KUPIEC TEST - MOVING AVERAGE APPROACH, 30 DAY LAG													
		95,00%				99,00%				99,90%			
Asset	No.	MIN	TARGET	RESULT	MAX	MIN	TARGET	RESULT	MAX	MIN	TARGET	RESULT	MAX
Brent Oil	3 259	**139**	**163**	**178**	**187**	**22**	**33**	**48**	**43**	**0**	**3**	**13**	**6**
OMXs30	3 215	**137**	**161**	**181**	**184**	**22**	**32**	**65**	**43**	**0**	**3**	**21**	**6**
STB3M	3 225	**137**	**161**	**170**	**185**	**22**	**32**	**96**	**43**	**0**	**3**	**47**	**6**

Table 4.1.2 Backtesting results with moving average approach.

The moving average approach underestimates the VaR at higher confidence levels due to the fatter tail of the returns than assumed by the normal distribution. The approach seams to perform well on lower confidence levels where the normality assumption is more valid. Consequently, the Kupiec test rejects the moving average approach at the higher confidence levels and accepts the approach at the 95% level.

As discussed in the method chapter, the characteristics of the OMXs30 returns give an indication that the assumption of normality is not that unreasonable. The OMXs30 returns are not that skewed and the kurtosis is not that great. Despite this, the approach performs better on the Brent oil returns even though the characteristics of those returns look to be further from the normality assumption. This can perhaps be explained by looking at the plot of the returns of the assets. The returns for Brent oil is more skewed and leptokurtic than those of the OMXs30, but the returns of the Brent oil are more stable whereas the OMXs30 returns show clear signs of volatility clustering.

The extreme volatility peaks of the STB3M returns makes for poor compatibility with the moving average approach at higher confidence levels, which can be seen in table 4.1.2. The approach produces approximately the double amount of VaR breaks, compared to the other two data sets. This is a result of the outliers often appear alone without a prior rise in volatility the days before the VaR break, as is usual when volatility is clustered. This is also the reason for the returns being so leptokurtic. This means that the approach has no way of forecasting these sudden rises in volatility and thus a VaR break occurs. The day after the VaR break when the outlier passes in to the measurement window the VaR measure rises significantly. Since the VaR break was a single occurrence it is not followed by additional extreme outcomes.

4.1.3 GARCH APPROACH

Table 4.1.3 shows the backtesting results of the GARCH approach.

KUPIEC TEST - GARCH														
		95,00%				99,00%				99,90%				
Asset	No.	MIN	TARGET	RESULT	MAX	MIN	TARGET	RESULT	MAX	MIN	TARGET	RESULT	MAX
Brent Oil	3 259	**139**	**163**	**156**	**187**	**22**	**33**	**43**	**43**	**0**	**3**	**16**	**6**
OMXs30	3 215	**137**	**161**	**161**	**184**	**22**	**32**	**47**	**43**	**0**	**3**	**17**	**6**
STB3M	3 225	**137**	**161**	**95**	**185**	**22**	**32**	**38**	**43**	**0**	**3**	**19**	**6**

Table 4.1.3 Backtesting results with GARCH approach.

The GARCH approach produces good results for the 95 percent confidence level with exception of the STB3M. The results for the 99% confidence level are also reasonable. The approach produces a bit too many VaR breaks as expected from a approach based on the normal distribution. The better performance of the GARCH approach on the 99% confidence level, compared to the moving average approach, is as we interpret it solely due to the GARCH approaches' more advanced way of estimating the daily variance. The GARCH approach does however also suffer from the same deficiency as the moving average approach which is that it assumes the returns to be normally distributed. This becomes apparent at the 99.9% confidence level, where the GARCH approach constantly produced to many VaR breaks to be accepted by the Kupiec test.

The Kupiec test rejects the approach for all of the assets at the 99.9% level. At the 99% level, the only asset for which the approach was rejected was the OMXs30. This was surprising at first, since the GARCH approach performed so well on the OMXs30 at the 95% confidence level. The reason to why the approach produced too many VaR breaks is probably that the assumption of normality is acceptable at lower confidence levels but not at higher ones. The further out in the tail of the distribution of returns we come, the less acceptable the assumption of normality becomes. This is however conflicted by the fact that the Kupiec test accepted the approach at the 99% confidence level for the Brent Oil and STB3M returns. This can be explained by the observation that the approach produces fewer VaR breaks than the target for these returns at the 95% level, i.e. the approach produces to high VaR values. At the 99% level it would seem that the assumption of normality that generally causes the approach to produce too many VaR breaks is countered by another error in the approach, which cancels both errors out, resulting in the approach being accepted by the Kupiec test.

The next question is to figure out what the error in the approach is that results in the approach producing too few VaR breaks at the 95% confidence level for Brent Oil and STB3M. We reason that it is yet again the assumption of normality that is throwing our calculations off. The maximum likelihood estimation formula that we used to estimate the values of ω, α and β also assumed that the returns that it was evaluating were normal. As argued in the method chapter, the returns of Brent oil and STB3M exhibited the characteristics least compatible with the normal distribution. We therefore suspect that the assumption of normality has caused the MLE function to produce values of ω, α and β that are slightly off their optimum, which in turn could be the cause of the GARCH function producing too few VaR breaks. This argument is supported by looking at the graphs of the values of ω, α and β in appendix 3 for all of our three assets. The graphs show how the values of ω, α and β change as we estimate them with a rolling window of the previous 2000 observations. The values for the OMXs30 returns are stable whereas the values for the two other assets fluctuate somewhat, indicating that the MLE function might not function properly with the returns as a result of the assumptions of normality.

These two errors both spring from the normality assumption, but seem to by chance affect the VaR in opposite direction, causing the approach to perform better than expected at the 99% confidence level. So even though the approach is accepted by the Kupiec test for Brent oil and STB30 at the 99% confidence level we have our reservations of whether the outcome really is a result of the approach performing well or if it is just a random occurrence.

Altogether, we feel that the GARCH approach does a good job of handling volatility clustering and estimating a variance forecasts. We do, however, feel that the MLE function that we have used does not work proficiently and that it would have to be adjusted to better fit the assets returns in order to further increase the performance of the GARCH approach.

According to Goorbergh & Vlaar the most important return characteristic to account for when calculating VaR is volatility clustering. We find that volatility clustering indeed is important but do not agree that it is the most important characteristic. As we have shown in this chapter the moving average approach performs well at the 95% confidence level even though it does not take volatility clustering into consideration. In our opinion the most important return characteristic is the distribution of the returns and how well they match the distribution assumed by the approach. Volatility clustering is according to our results the second most important characteristic.

4.2 EXACTNESS & SIMPLICITY

The tradeoff between exactness and simplicity is an interesting but hard dilemma to solve or to decide about. People will argue that if a model do not provide exact and accurate answers or estimations, they are not useful. Some models have to provide results or numbers that is to a 100 percent correct, otherwise it can result in severe outcomes. A straight forward example is models used in the construction of airplanes, of which its functions is directly related to the outcome of the specific models. The models or approaches used to calculate VaR is estimates of future returns. When estimating future events one can never be a 100 percent sure about the outcome. This leads to approaches trying to come as close as possible to a likely future. By constructing more or less advanced approaches the exactness is hoped to be found. An approach based on many parameters can be useful, however if the parameters put into the approach is not correct the complexity will not be to any benefit of the approach and will only result in bad outcomes. This study has examined both a non parametric approach and two more complex parametric approaches. The biggest and most important assumption that the parametric approaches takes is the assumption of normal distribution of returns. When examining the characteristics of the asset returns, it becomes clear that the normal distribution provides a poor fit. This affects both the parametric approaches, moving average and GARCH negative and results in both underestimations and overestimations of the VaR.

From the backtesting results it becomes clear that the approach performing best on the two lowest confidence levels 95 percent and 99 percent is the GARCH approach and the best performance on the 99.9 percent level comes from the historical simulation. The parametric approach outperforms the non parametric approach on the levels where the normal distribution has a smaller effect on the result. It does however effect the result, which leads to the fact that volatility clustering is an important parameter. The moving average approach that uses normal distribution but disregards volatility clustering is the approach that performs worst. When the highest confidence level is used, the effect of volatility clustering used in GARCH cannot overtake the fact of the normal distribution assumption, leading to best performance coming from the historical simulation.

The tradeoff complexity of simplicity versus exactness does not have a single answer. It depends. It is obvious than one wants exactness with as much simplicity as possible. The parametric approaches are more complex to use and work with. They are a lot more time-consuming than the historical simulation, leading to higher implementation costs. In many cases, the normal distribution, which is the distribution often used in parametric approaches, provides a bad fit. It is, however, widely used because of the many benefits that come with it. If the desired confidence level is high, it is important to be aware of the negative effects the assumption of normality will have on the results. If the effect of volatility clustering will cause no big improvements to the approach, the historical simulation may be a better option. If the normal distribution on the other hand provides a good fit for the asset returns the parametric approach and more specific the GARCH approach seems to give superior results. This leads us to believe that the characteristics of the returns are highly correlated to the results of the different approaches and confidence levels chosen.

4.3 VALIDITY

The validity of a study is of great importance. When it comes to the surface validity, we can conclude that the results relates to earlier studies, even though earlier studies perhaps have not been conducted on the same conditions using the same assets we based our study on. But the fact that the normal distribution used in the parametric approaches did not suit the asset returns very well can be seen in other results as well. The same is true about the usefulness of the historical simulation approach. It has been stated before that the non parametric approach is rather simple but performs relatively good compared to parametric approaches. It even outperforms the other approaches on the higher levels of confidence. This is because the returns fit badly with the normal distribution.

As we have discussed previously the results from earlier studies have been somewhat conflicting. One reason for this might be that the different results from these studies are hard to compare to one another since they were conducted on different assets at different points in time. The approaches are also almost never applied in the same way, they are often adjusted in some way to better fit the asset examined. In accordance with Linsmeier & Pearson we find that there is no simple answer to which VaR approach is the best. Instead it would seem that the approaches perform differently under various circumstances with different assets. Therefore one has to evaluate which approach is best suited for each individual asset and situation.

If we look closer on the internal validity, the results are about what we expected, however with some divergences. Again concentrating on the histograms presented in chapter three and the characteristics of the asset returns presented in the same, we found an indication of a poor fit between the asset returns and the normal distribution. We did however assume that the results of the OMXs30 under the parametric approach, and then especially for the GARCH approach, would be better than for the Brent oil since the characteristics of OMXs30 more assembled the distribution of normality. This was not the case except for the 95 percent confidence level. The worst results under all approaches was for the STB3M, a result that we anticipated but perhaps believed that the historical simulation would handle a bit better than it did. The truth is that the non parametric approach could have produced a better result if we had chosen to shorten the amount of days that the volatility was based on. By doing this, the extreme daily volatility in the beginning of the 1990s would have had less impact of the result as the historical simulation is an approach that responds very slowly to changes in volatility.

The degree of external validity can be seen as relatively high even though the results indicates that there is no single truth about the usefulness of the VaR approaches. It all depends on what one wants to get out of the VaR measure. The results all show on the non perfect fit between normal distribution and asset returns, leading to the necessity of establishing the distribution that fits the returns best. The characteristics of the returns are of highest importance, and when the distribution is established, the results show that the GARCH approach would produce relatively accurate values and most certainly outperform the historical simulation also on the higher levels of confidence.

5. CONCLUSIONS

This chapter accounts for the conclusions that the study has reached and suggests topics for further research.

5.1 CONCLUSIONS

Table 5.1 sums up the results of the Kupiec test that we have conducted and based our analysis on.

REJECTION TABLE KUPIEC TEST				
	Underlying asset	95,00%	99,00%	99,90%
HISTORICAL SIMULATION	*Brent Oil*	REJECTED	ACCEPTED	REJECTED
	OMXs30	REJECTED	REJECTED	ACCEPTED
	STB3M	REJECTED	ACCEPTED	REJECTED
MOVING AVERAGE	*Brent Oil*	ACCEPTED	REJECTED	REJECTED
	OMXs30	ACCEPTED	REJECTED	REJECTED
	STB3M	ACCEPTED	REJECTED	REJECTED
GARCH	*Brent Oil*	ACCEPTED	ACCEPTED	REJECTED
	OMXs30	ACCEPTED	REJECTED	REJECTED
	STB3M	REJECTED	ACCEPTED	REJECTED

Table 5.1 Summary of backtesting results.

As the table shows, no approach was accepted for all the levels of confidence for all the assets. The test does however indicate that the nonparametric approach performs better at higher levels of confidence and the parametric approaches perform better at lower levels. The conclusion is that the assumption of normally distributed returns is unacceptable and affects the results negatively for the parametric approaches at the 99% and 99.9% confidence levels. It could however be a reasonable approximation at the 95% confidence level for some assets with suitable returns. Another conclusion that we have drawn is that there are very few general truths when it comes to time series analysis, each asset has to be studied individually to evaluate which approach to use under different circumstances and levels of confidence.

Looking back at our calculations we feel that it would have been appropriate to set different window sizes for each of the different confidence levels under the historical simulation approach. A window size of 2000 observations was too large for the 95% confidence interval. There seems to be a relationship between the kurtosis of the returns, the confidence level and the appropriate historical window size. The level of confidence gives an indication of what could be an appropriate window size, and the window size must then be adjusted to fit the kurtosis of the returns in order to capture the right amount of extreme outliers that corresponds to the confidence level.

The trade-off between exactness and simplicity is very interesting. Many econometricians would find the GARCH approach used in this study quite basic and straight forward. For someone not familiar with time series analysis the approach might seem too complex to implement, especially with regard to the maximum likelihood estimation. Looking at the result the GARCH approach does not perform as well as would be expected in relation to its complexity. When comparing it to the historical and moving average approach, which can easily be implemented using basic tools such as Excel, the GARCH approach seems like an unnecessary effort given the little extra exactness added. It should however be noted that there are many different ways in which these approaches can be implemented and it might very well be that the GARCH approach could perform considerably better if implemented in another way than the one tested here. The moving average approach does however perform surprisingly well at the 95% confidence level, and we feel that for small scale use where maximum precision is not crucial, this approach is adequate at low confidence levels. The moving average approach could then be complemented with historical simulation for use at higher confidence levels to provide a simple and effective way to estimate VaR for all confidence levels.

We would like to end this paper by answering the questions stated in the problem discussion.

- How do the different characteristics of the returns of the underlying assets affect VaR?

Assets whose return characteristics does not stray too far from the assumption of normality can satisfactorily be evaluated with the parametric approaches used in this study at lower levels of confidence. At higher levels of confidence, the historical simulation approach is more appropriate. Even though the GARCH approach evaluated in this study compensates for volatility clustering it still suffers from the same flaw as the moving average approach at the higher confidence levels. It produces good results at the 95% confidence level for assets where the assumption of normality is not that farfetched.

- Which of the approaches gives the most accurate VaR measure?

A conclusive answer to this question cannot be found using the results in this study. The conclusion is that the asset return characteristics determine which approach that performs best at different levels of confidence. A conscious evaluation of asset returns and other case specific factors should be made in each case where VaR is to be implemented, in order to choose the right approach.

5.2 FURTHER RESEARCH

This study has been conducted on three assets using three different approaches and three different levels of confidence. Therefore are the results based on the specific characteristics of the returns of the assets chosen even though the conclusion drawn is that the results can be used in other studies.

It would be interesting to conduct a study on other assets to see if the conclusions drawn would be the same. Other assets can have characteristics that fit the normal distribution better leading to better results from the parametric approaches. It would further be interesting to base the parametric approaches on other distributions such as for example the t-distribution that better handles the fact of fat tails. Using more types of approaches could further enhance the study.

A subject for further research could be to examine suitable window sizes for the calculation of historical simulation VaR. The window size is the only factor that influences the VaR value which makes it crucial to the approach.

The largest source of error that we have discovered in this study is the assumption made by the parametric approaches that the returns are normally distributed. Further research in simple VaR approaches such as the moving average approaches which utilize alternative probability distributions would be interesting.

REFERENCES

LITERARY SOURCES

Arbnor, I., Bjerkner, B., 1994. *Företagsekonomisk metodlära*. 2nd ed. Studentlitteratur

Cabedo, J. D., Moya, I., 2003. Oil price 'Value at Risk' using the historical simulation approach. *Energy Economics*, 2, pp 239–253.

Costello, A., Asem, E., Gardner, E., 2008. Comparison of historically simulated VaR: Evidence from oil prices. *Energy Economics,* 30, pp 2154–2166.

Dowd, K., 1998. *Beyond value at risk: The new science of risk management*. John Wiley & Sons Ltd.

Eydeland, A., Wolyniec, K., 2003. *Energy and Power Risk Management, New Developments in Modeling, Pricing and Hedging*. John Wiley & Sons, Inc.

Fernandez, V., 2003. Extreme value theory and value at risk. *Revista de Analisis Economic*, Vol 18 (1), pp 57-85.

Goorbergh, R., Vlaar, P., 1999. Value-at-Risk Analysis of Stock Returns, Historical Simulation, Variance Techniques or Tail Index Estimation? *Econpapers, Working Papers Series*, Netherlands Centralbank.

Hendricks, D., 1996. Evaluation of value-at-risk models using historical data, *Economic Policy Review*, April, Vol. 2.

Holme, I. M., Solvang, B. K., 1991. *Forskningsmetodik – Om kvalitativa och kvantitativa metoder*. Studentlitteratur.

Holton, G. A., 2002. History of Value-at-Risk: 1922-1998. *Econpapers, Working paper,* July 25.

Hull, J. C., 2008. *Options futures and other derivatives*. 7th ed. Prentice Hall.

Jorion, P., 2001. *Value at risk*. 2nd Ed. McGraw-Hill.

Jorian, P., 2002. How Informative Are Value-at-Risk Disclosures? *The Accounting Review*, Vol. 77, No. 4 pp. 911-931.

Lee, C. F., Lee, J. C., & Lee, A. A., 2000. *Statistics for business and financial economics*. 2nd ed. World Scientific.

Linsmeier, T. J., Pearson, N. D., 1996. Risk measurement: An introduction to Value at Risk. *Econpapers, Working Paper*.

Mauro, A., 1999. Price Risk Management in the Energy Industry: The Value at Risk Approach. *Social Science Research Network.*

Penza, P., Bansal, V. K., 2001. *Measuring market risk with value at risk*. John Wiley & Sons.

Sadeghi, M., Shavvalpour, S., 2006. Energy risk management and value at risk modeling. *Energy Policy,* 34, pp 3367–3373.

Sadorsky, P., 1999. Oil price shocks and stock market activity. *Energy Economic,* 21, pp 449–469.

Saunders, A., 1999. *Credit risk management*. John Wiley& Sons.

Saunders, A., Cornett, M. M., 2007. *Financial Markets and Institutions*. McGraw-Hill Irwin.

Tsay, R. S., 2002. *Analysis of Financial Time Series*, John Wiley& Sons.

Veiga, Chan & McAleer, 2008. Evaluating the impact of market reforms on Value-at-Risk forecasts of Chinese A and B shares. *Pacific-Basin Finance Journal*, Sep, Vol. 16/4, pp 453-475.

INTERNET SOURCES

http://tonto.eia.doe.gov/dnav/pet/hist/rbrted.htm (2008-10-13)

http://support.microsoft.com/kb/214115 (2008-12-15)

http://www.omxnordicexchange.com/produkter/index/OMX_index/OMXS_Local_Index/
(2008-11-14)

http://www.riksbank.se/templates/stat.aspx?id=16739 (2008-11-13)

APPENDIX

The following pages contain the appendixes of the essay.

APPENDIX 1

In the text box below the macro is presented that was used to solve the maximum likelihood estimation in Excel.

```
Sub MLE()
Do While ActiveCell.Offset(100, 0) <> ""
      SolverReset
      SolverOptions MaxTime:=300
      SolverOKSetCell:=ActiveCell.Offset(0, 0), MaxMinVal:=1, ValueOf:="0",
ByChange:=Range("$F$5471:$H$5471")
      SolverAdd CellRef:="$F$5471", Relation:=1, FormulaText:="$F$5474"
      SolverAdd CellRef:="$F$5471", Relation:=3, FormulaText:="$F$5475"
      SolverAdd CellRef:="$G$5471", Relation:=1, FormulaText:="$F$5474"
      SolverAdd CellRef:="$G$5471", Relation:=3, FormulaText:="$F$5475"
      SolverAdd CellRef:="$H$5471", Relation:=1, FormulaText:="$F$5474"
      SolverAdd CellRef:="$H$5471", Relation:=3, FormulaText:="$F$5475"
      SolverAdd CellRef:="$I$5471", Relation:=1, FormulaText:="$F$5474"
      SolverSolve UserFinish:=True
 With ActiveCell.Interior
         .Pattern = xlSolid
         .PatternColorIndex = xlAutomatic
         .Color = 5296274
         .TintAndShade = 0
        .PatternTintAndShade = 0
   End With
   Range("$F$5471").Copy Destination:=ActiveCell.Offset(0, 1)
   Range("$G$5471").Copy Destination:=ActiveCell.Offset(0, 2)
   Range("$H$5471").Copy Destination:=ActiveCell.Offset(0, 3)
   Application.Wait (Now + TimeValue("0:00:03"))
   ActiveCell.Offset(100, 0).Activate
   Loop
End Sub
```

APPENDIX 2

Tests for autocorrelation in the residuals of a lagged regression of the returns of the assets.

BRENT OIL

Residuals:

Min	1Q	Median	3Q	Max
-0.3627667	-0.0116968	0.0002715	0.0122172	0.1786131

Coefficients:

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	0.0002045	0.0003142	0.651	0.51521
Brent Oil	0.0389464	0.0135341	2.878	0.00402 **

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 0.0232 on 5451 degrees of freedom

Multiple R-squared: 0.001517

OMXs30

Residuals:

Min	1Q	Median	3Q	Max
-0.0854533	-0.0075695	0.0004188	0.0078160	0.1083411

Coefficients:

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	0.0002633	0.0001993	1.321	0.18643
OMXs30	0.0404256	0.0135128	2.992	0.00279 **

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 0.01474 on 5472 degrees of freedom
Multiple R-squared: 0.001633
Adjusted R-squared: 0.00145
F-statistic: 8.95 on 1 and 5472 DF
p-value: 0.002787

STB3M

```
Residuals:
       Min          1Q      Median        3Q        Max
-0.2875080  -0.0030782  0.0001740  0.0028235  0.3307909

Coefficients:
              Estimate    Std. Error   t value    Pr(>|t|)
(Intercept)  -0.0001740   0.0001988    -0.876     0.381
ssvx$ktwo     0.0584564   0.0135053    4.328      1.53e-05 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.0147 on 5465 degrees of freedom
Multiple R-squared: 0.003416
Adjusted R-squared: 0.003234
F-statistic: 18.74 on 1 and 5465 DF
p-value: 1.529e-05
```

APPENDIX 3

The following graphs show the values of α, β and ω and how they changed during time when the MLE was applied to a window of the last 2000 observations.

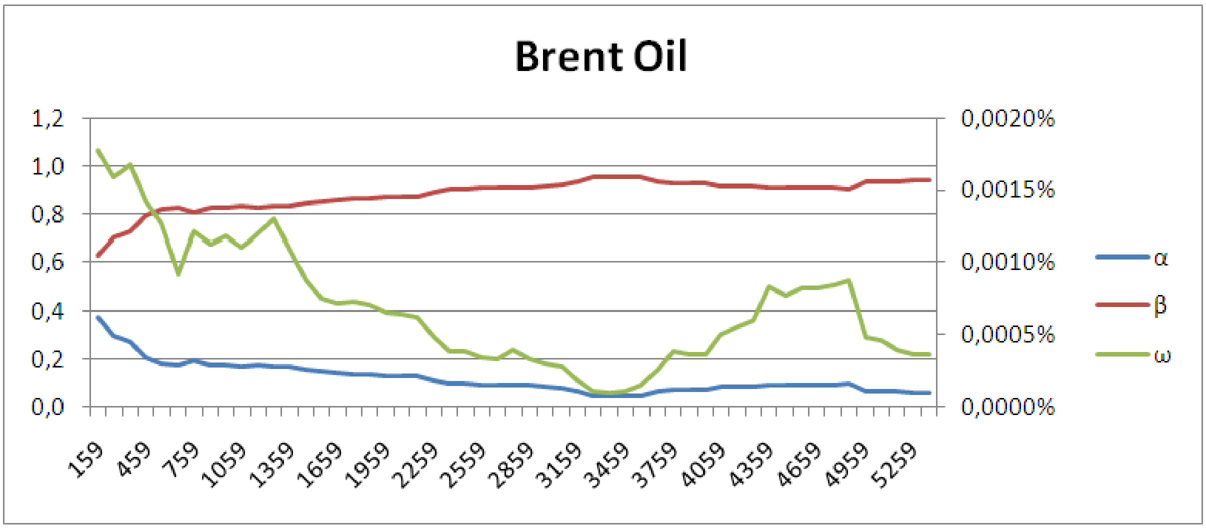

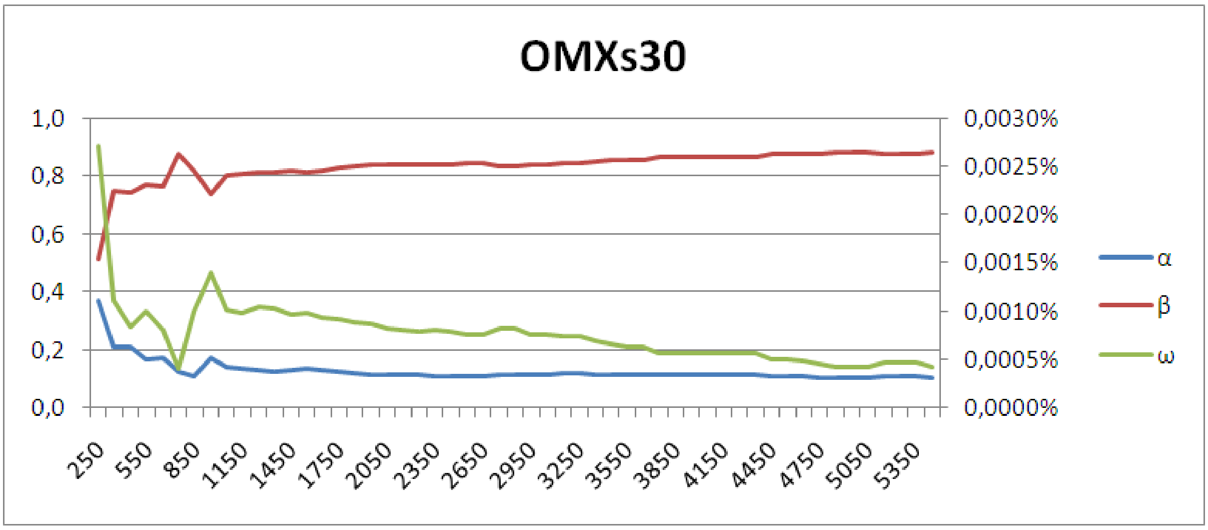

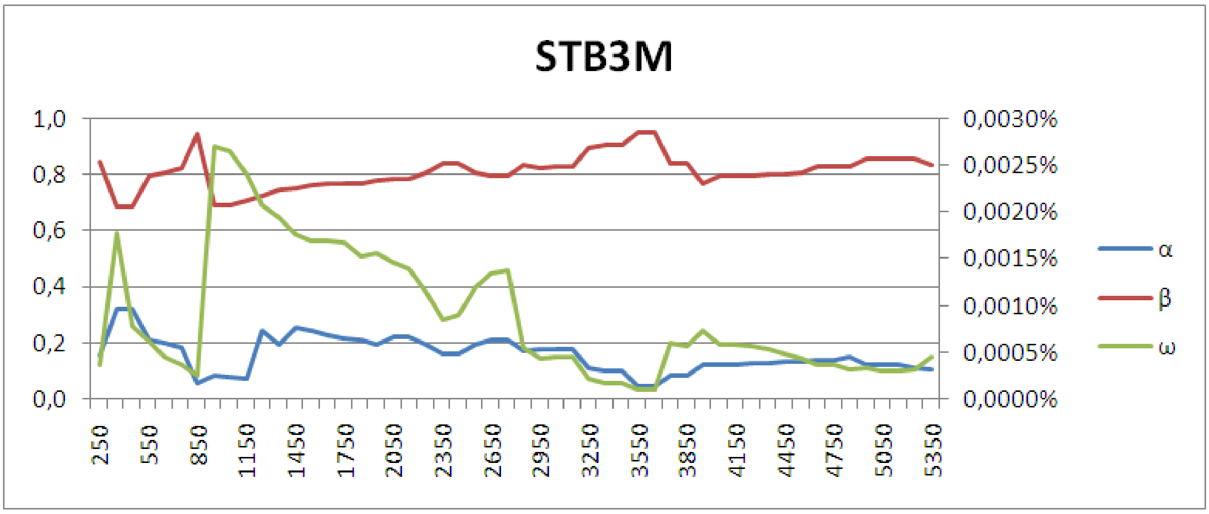

APPENDIX 4

In this appendix all the graphs for the different VaR measures for the different confidence levels will be listed.

HISTORICAL SIMULATION APPROACH

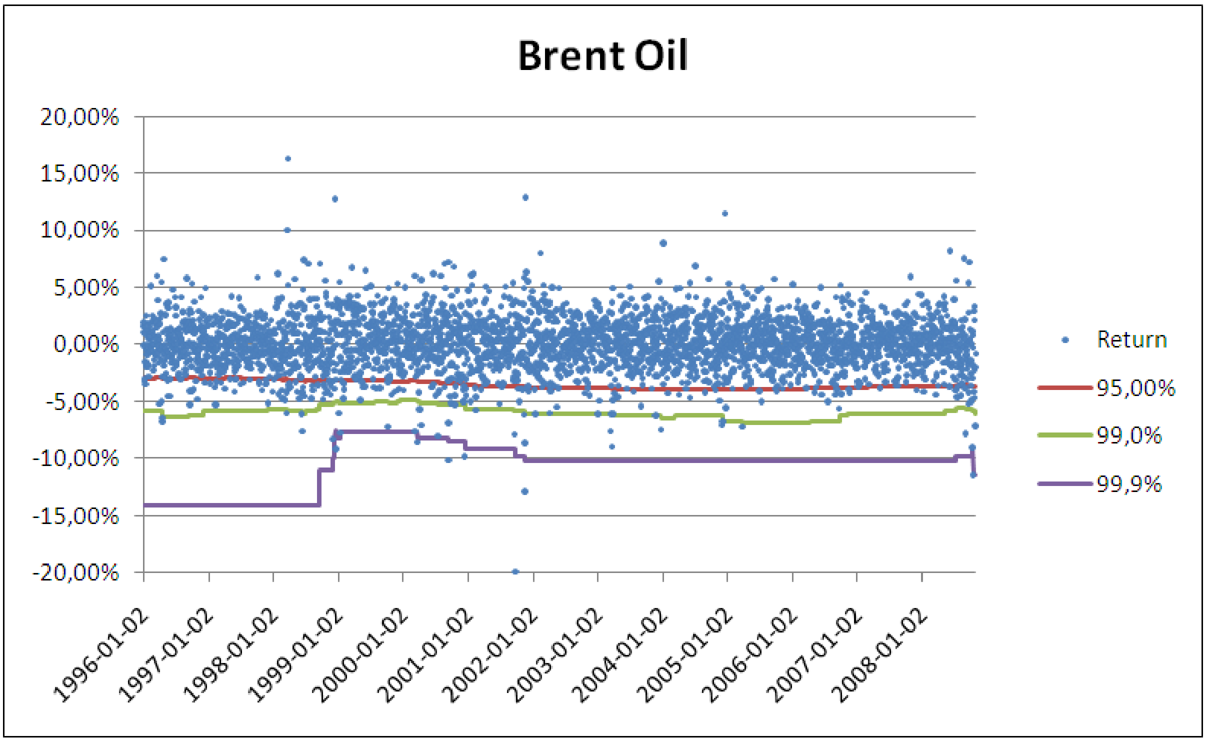

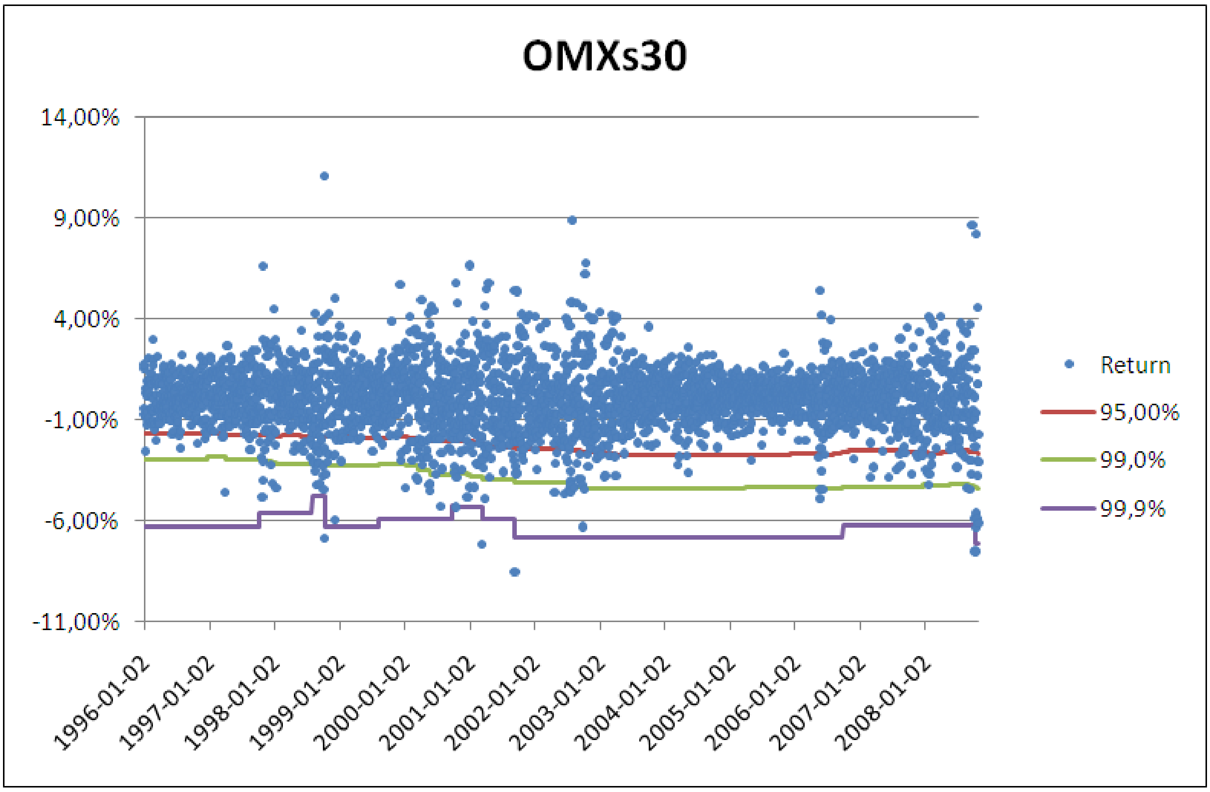

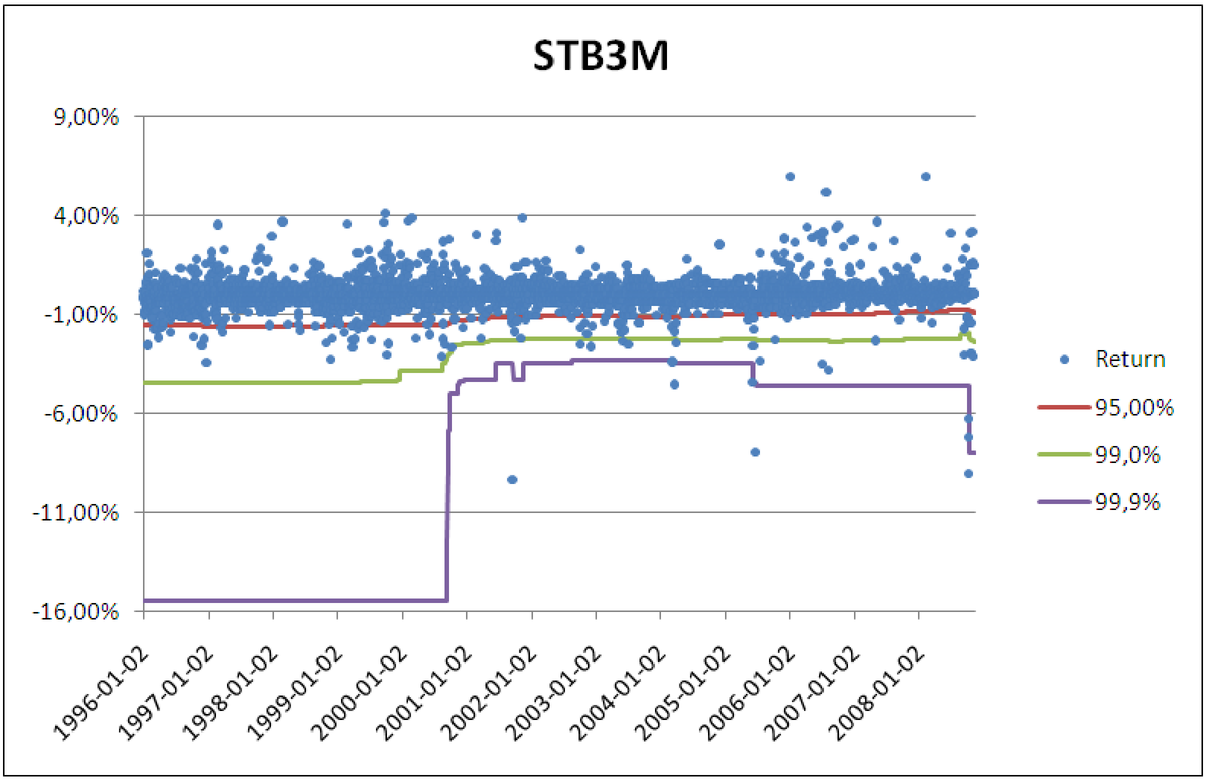
STB3M
9,00%
4,00%
-1,00%
-6,00%
-11,00%
-16,00%
1996-01-02
1997-01-02
1998-01-02
1999-01-02
2000-01-02
2001-01-02
2002-01-02
2003-01-02
2004-01-02
2005-01-02
2006-01-02
2007-01-02
2008-01-02
Return
95,00%
99,0%
99,9%

MOVING AVERAGE APPROACH

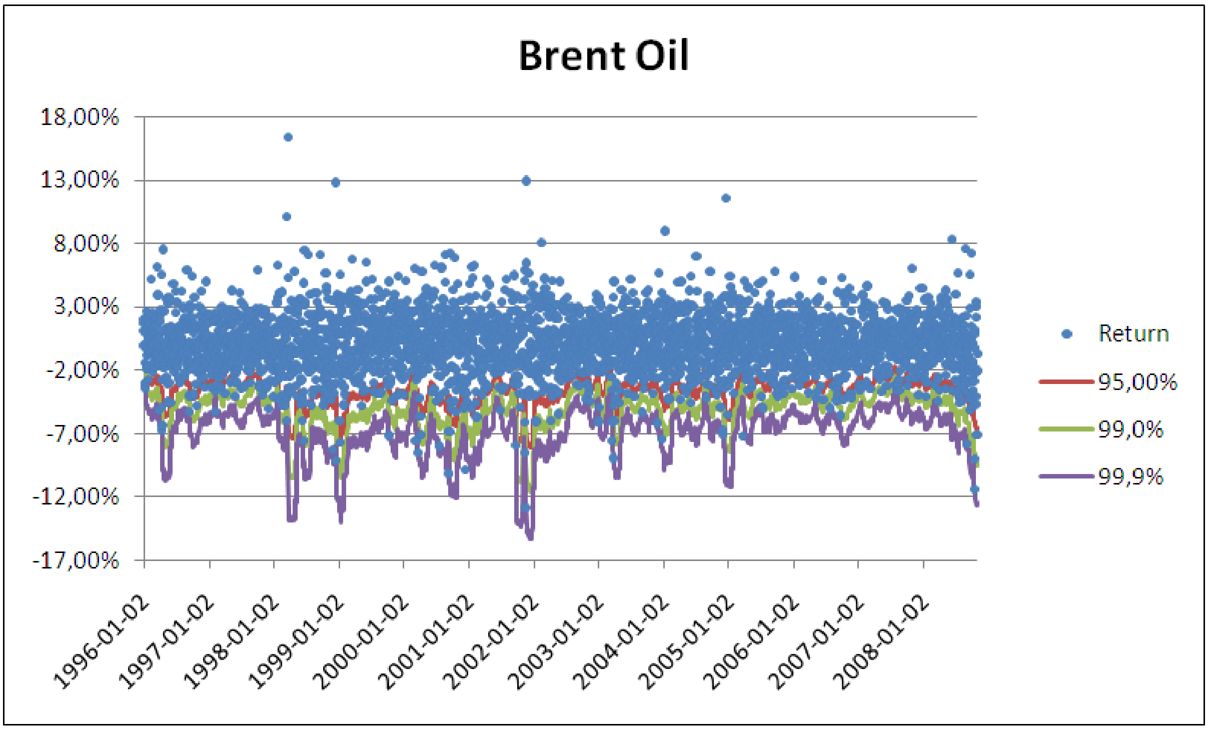

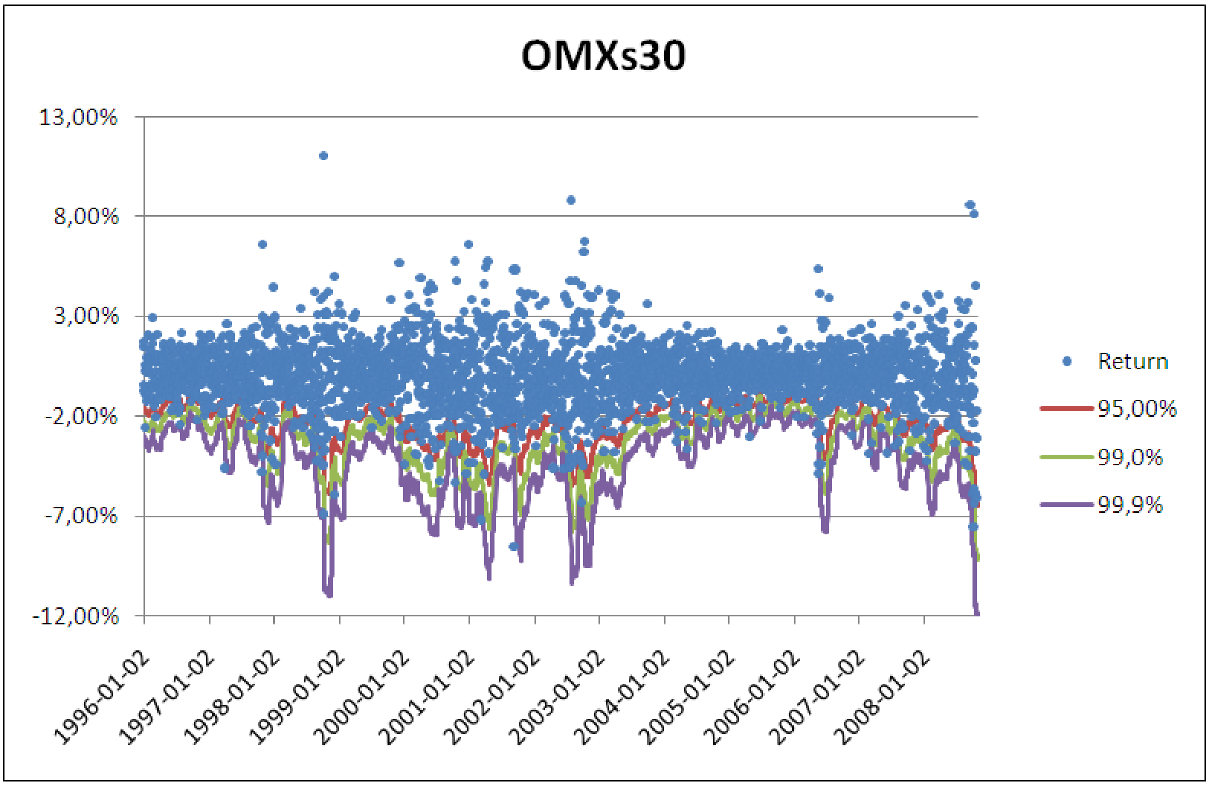

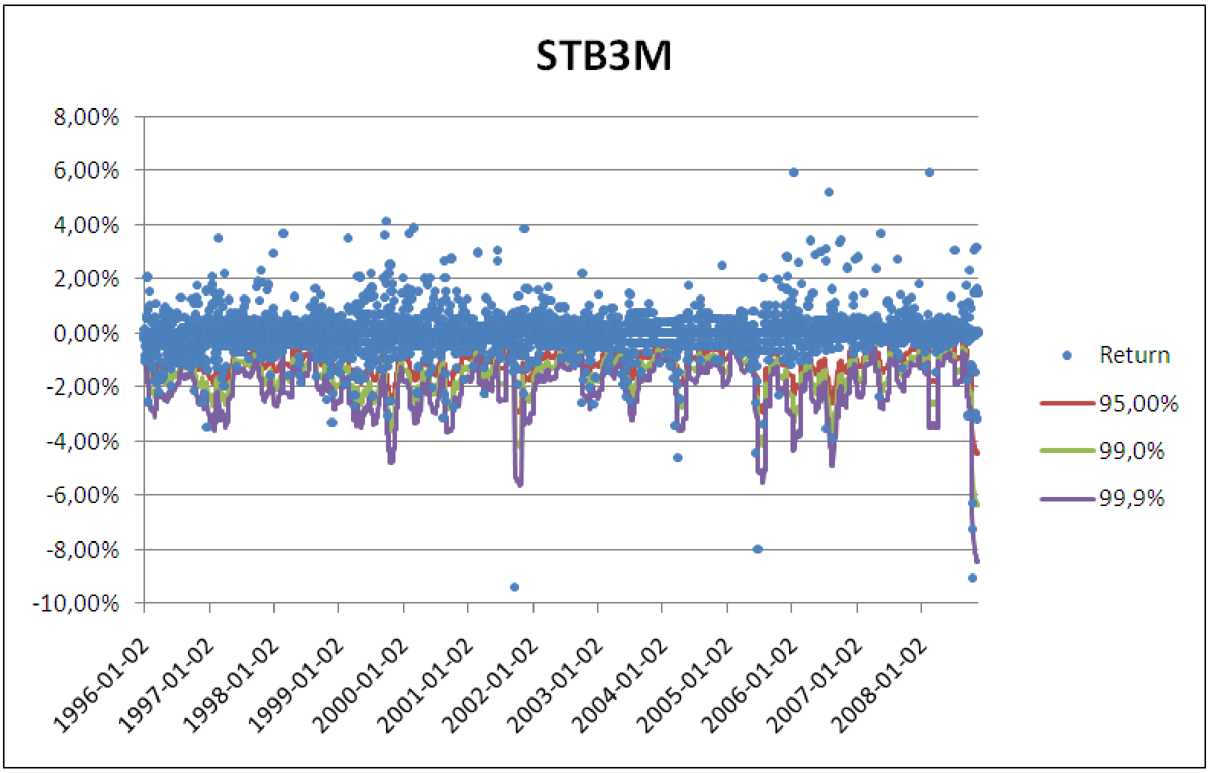
STB3M
8,00%
6,00%
4,00%
2,00%
0,00%
-2,00%
-4,00%
-6,00%
-8,00%
-10,00%
1996-01-02
1997-01-02
1998-01-02
1999-01-02
2000-01-02
2001-01-02
2002-01-02
2003-01-02
2004-01-02
2005-01-02
2006-01-02
2007-01-02
2008-01-02
Return
95,00%
99,0%
99,9%

GARCH APPROACH

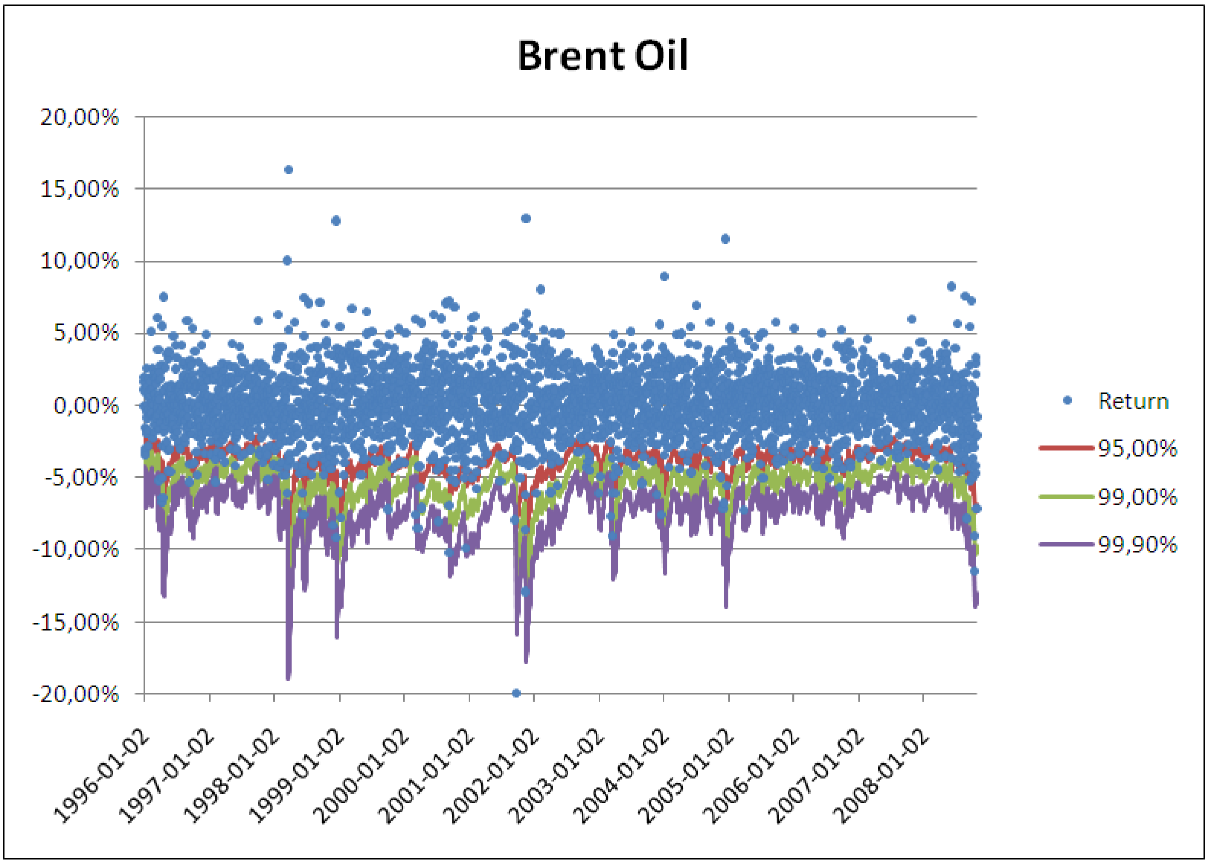

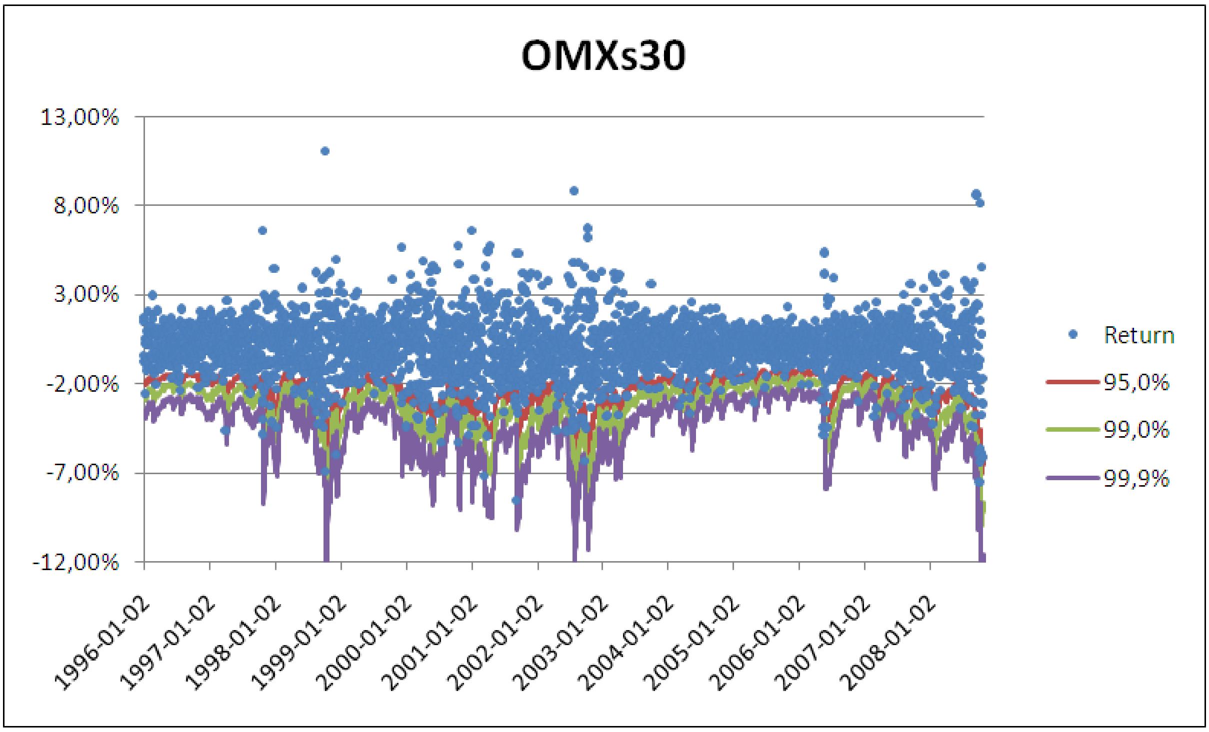

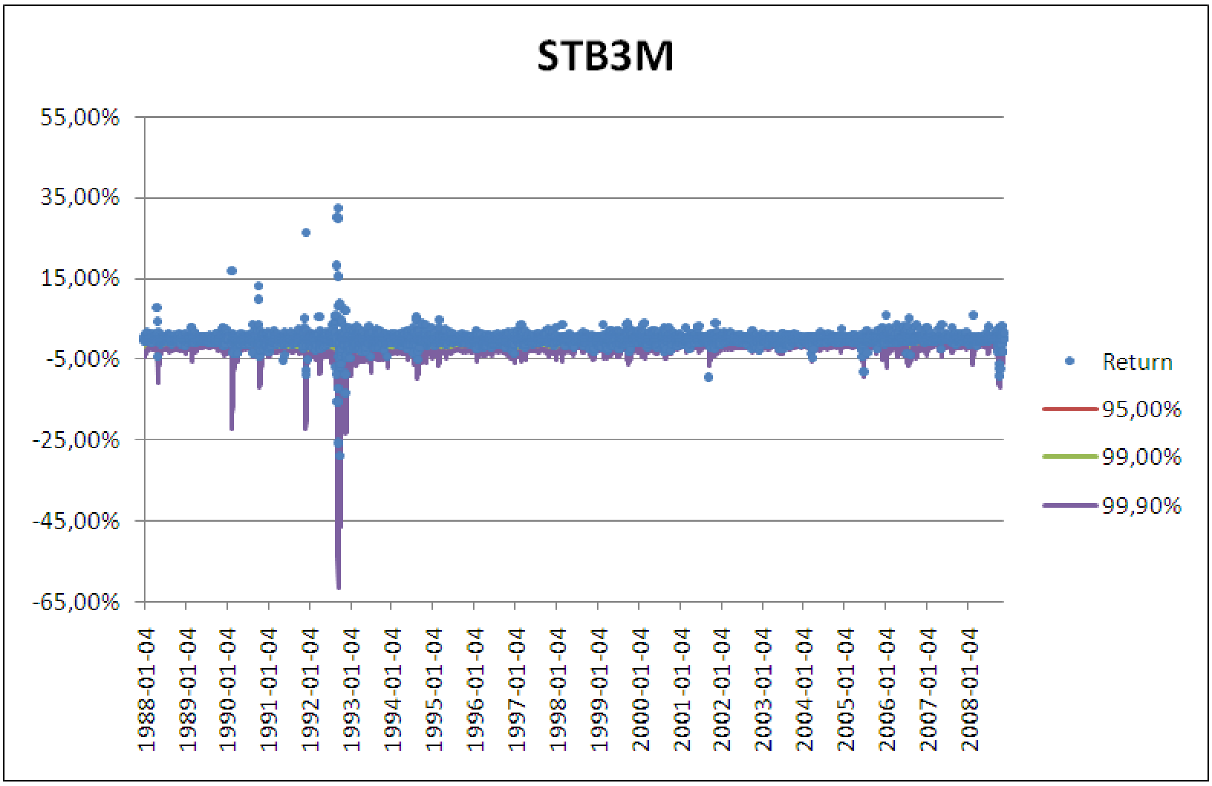
STB3M
55,00%
35,00%
15,00%
-5,00%
-25,00%
-45,00%
-65,00%
1988-01-04
1989-01-04
1990-01-04
1991-01-04
1992-01-04
1993-01-04
1994-01-04
1995-01-04
1996-01-04
1997-01-04
1998-01-04
1999-01-04
2000-01-04
2001-01-04
2002-01-04
2003-01-04
2004-01-04
2005-01-04
2006-01-04
2007-01-04
2008-01-04
Return
95,00%
99,00%
99,90%